Phase Transitions of Spectral Initialization for High-Dimensional Nonconvex Estimation

Yue M. Lu and Gen Li

arXiv:1702.06435v2 [cs.IT] 11 Apr 2017

***Abstract*—We study a spectral initialization method that serves a key role in recent work on estimating signals in nonconvex settings. Previous analysis of this method focuses on the phase retrieval problem and provides only performance bounds. In this paper, we consider arbitrary generalized linear sensing models and present a precise asymptotic characterization of the performance of the method in the high-dimensional limit. Our analysis also reveals a phase transition phenomenon that depends on the ratio between the number of samples and the signal dimension. When the ratio is below a minimum threshold, the estimates given by the spectral method are no better than random guesses drawn from a uniform distribution on the hypersphere, thus carrying no information; above a maximum threshold, the estimates become increasingly aligned with the target signal. The computational complexity of the method, as measured by the spectral gap, is also markedly different in the two phases. Worked examples and numerical results are provided to illustrate and verify the analytical predictions. In particular, simulations show that our asymptotic formulas provide accurate predictions for the actual performance of the spectral method even at moderate signal dimensions.**

***Index Terms*—Spectral initialization, signal estimation, nonconvex optimization, spiked covariance model, phase transition**

I. Introduction

We consider the problem of estimating an n-dimensional vector $\boldsymbol{\xi}$ from a number of generalized linear measurements. Let $\{\boldsymbol{a}_i\}_{1\le i\le m}$ be a set of sensing vectors in $\mathbb{R}^n$. Given $\{\boldsymbol{a}_i^T\boldsymbol{\xi}\}$, the measurements are drawn independently from

$$y_i \sim f(y\,|\,\boldsymbol{a}_i^T\boldsymbol{\xi}), \tag{1}$$

where $f(\cdot\,|\,\cdot)$ is a conditional density function modeling the acquisition process. This model arises in many problems in signal processing and statistical learning. Examples include photon-limited imaging [1], [2], phase retrieval [3], signal recovery from quantized measurements [4], and various single-index and generalized linear regression problems [5], [6].

The standard method for recovering $\boldsymbol{\xi}$ is to use the estimator

$$\widehat{\boldsymbol{\xi}} = \arg\min_{\boldsymbol{x}} \sum_{i=1}^{m} \ell(y_i, \boldsymbol{a}_i^T\boldsymbol{x}), \tag{2}$$

where $\ell : \mathbb{R}^2 \to \mathbb{R}$ is some loss function (*e.g.*, the negative log-likelihood of the observation model as used in maximum likelihood estimation.) In many applications, however, the natural loss function is not convex with respect to $\boldsymbol{x}$. There is often no effective way to convexify (2). In those cases for which convex relaxations do exist, the resulting algorithms can be computationally expensive. The problem of phase retrieval, where $y_i = (\boldsymbol{a}_i^T\boldsymbol{\xi})^2 + \varepsilon_i$ for some noise terms $\{\varepsilon_i\}$, is an example in the latter scenario. Convex relaxation schemes such as those based on lifting and semidefinite programming (*e.g.*, [7]–[10]) have been successfully developed for solving the phase retrieval problem, but the challenges facing these schemes lie in their actual implementation. In practice, the computational complexity and memory requirement associated with these convex-relaxation methods are prohibitive for signal dimensions that are encountered in real-word applications such as imaging.

In light of these issues, there is strong recent interest in developing and analyzing efficient iterative methods that directly solve nonconvex forms of (2). Examples include the alternating minimization scheme for phase retrieval [11], the Wirtinger Flow algorithm and its variants [12]–[16], iterative projection methods [17], [18], and recent schemes for phase retrieval using linear programming [19], [20]. A common ingredient that contributes to the success of these algorithms for nonconvex estimation is that they all use some carefully-designed spectral method as an initialization step, which is then followed by further (iterative) refinement. Beyond the signal estimation problem considered in this paper, related spectral methods have also been successfully applied to initialize algorithms for solving other nonconvex problems such as matrix completion [21], low-rank matrix recovery [22], blind deconvolution [23], [24], sparse coding [25], and joint alignment from pairwise differences [26].

In this paper, we present an *exact* high-dimensional analysis of a widely-used spectral method [11]–[13] for estimating $\boldsymbol{\xi}$. The method consists of only two steps: First, construct a data matrix from the sensing vectors and measurements as

$$\boldsymbol{D}_m \overset{\text{def}}{=} \frac{1}{m}\sum_{i=1}^{m} \mathcal{T}(y_i)\boldsymbol{a}_i\boldsymbol{a}_i^T, \tag{3}$$

where $\mathcal{T} : \mathbb{R} \to \mathbb{R}$ is a preprocessing function (*e.g.* a trimming or truncation step.) Second, compute a normalized eigenvector, denoted by $\boldsymbol{x}_1$, that corresponds to the largest eigenvalue of $\boldsymbol{D}_m$. The vector $\boldsymbol{x}_1$ is then our estimate of $\boldsymbol{\xi}$ (up to an unknown scalar.) It is notable that this method is *model-free* in that the algorithm does not require the knowledge of the exact acquisition process [*i.e.*, the conditional density $f(\cdot\,|\,\cdot)$ in (1).]

Y. M. Lu is with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA (e-mail: yuelu@seas.harvard.edu). Part of this work was done during his visit to the Information Initiative at Duke (iiD) in Spring 2016. He thanks the members of this interdisciplinary program for their hospitality.

G. Li is with the Department of Electronic Engineering, Tsinghua University, Beijing 100084, China (e-mail: g-li16@mails.tsinghua.edu.cn). He was a summer visiting undergraduate student at the John A. Paulson School of Engineering and Applied Sciences, Harvard University.

This work was supported in part by the ARO under contract W911NF-16-1-0265 and by the US National Science Foundation under grant CCF-1319140. Preliminary version of this work will be presented at the IEEE International Symposium on Information Theory (ISIT) in 2017.

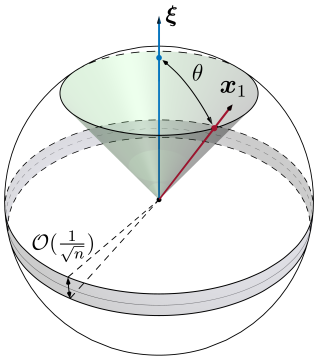

Fig. 1. Illustrations of the phase transitions of the spectral method. Depending on the sampling ratio α, the asymptotic performance of the spectral method can be in one of two very different phases: In the uncorrelated phase, its estimate $\boldsymbol{x}_1$ is asymptotically orthogonal to $\boldsymbol{\xi}$. The performance in this case is no better than an arbitrary guess drawn uniformly at random from the hypersphere $\mathcal{S}^{n-1}$. In the correlated phase, the estimate $\boldsymbol{x}_1$ (or its negative version $-\boldsymbol{x}_1$) will be concentrated on the surface of a right circular cone, making an angle $\theta = \arccos\left(\sqrt{\rho(\alpha)}\right)$ with the target vector $\boldsymbol{\xi}$.

This spectral method was introduced by Netrapalli, Jain, and Sanghavi in [11] to address the problem of phase retrieval. Under the assumption that the sensing vectors consist of i.i.d. Gaussian random variables, the authors show that the leading eigenvector $\boldsymbol{x}_1$ is aligned with the target vector $\boldsymbol{\xi}$ in direction when there are sufficiently many measurements. More specifically, they show that the squared cosine similarity

$$\rho(\boldsymbol{\xi}, \boldsymbol{x}_1) \overset{\text{def}}{=} \frac{(\boldsymbol{\xi}^T \boldsymbol{x}_1)^2}{\|\boldsymbol{\xi}\|^2 \|\boldsymbol{x}_1\|^2}, \tag{4}$$

which measures the degree of the alignment between the two vectors, approaches 1 with high probability, when the number of samples $m \geq c_1 n \log^3 n$. This sufficient condition on sample complexity was later improved to $m \geq c_2 n \log n$ in [7], and further improved to $m \geq c_3 n$ in [13] with an additional trimming step on the measurements. In these expressions, c_1, c_2, c_3 stand for some unspecified numerical constants.

In this paper, we provide a precise asymptotic characterization of the performance of the spectral method under Gaussian measurements. Although it was originally proposed for the special case of phase retrieval, the spectral method is applicable to a much broader class of models. Therefore, in our analysis, we consider general acquisition models under arbitrary conditional distributions $f(y \,|\, \boldsymbol{a}_i^T \boldsymbol{\xi})$. Moreover, unlike previous work, which only provides bounds for $\rho(\boldsymbol{\xi}, \boldsymbol{x}_1)$, we derive the exact high-dimensional limit of this value. In particular, we show that, as n and m both tend to infinity with the *sampling ratio* $\alpha \overset{\text{def}}{=} m/n$ kept fixed, the squared cosine similarity ρ converges in probability to a limit value $\rho(\alpha)$. Explicit formulas are provided for computing $\rho(\alpha)$.

Geometrically, the squared cosine similarity $\rho(\boldsymbol{\xi}, \boldsymbol{x}_1)$ as defined in (4) specifies the angle θ between $\boldsymbol{\xi}$ and $\boldsymbol{x}_1$. The values of ρ vary from 0 to 1: $\rho = 1$ means perfect alignment, *i.e.*, $\theta = 0$ or π; and $\rho = 0$ is the opposite case, meaning $\boldsymbol{x}_1$ is orthogonal to (*i.e.* uncorrelated with) $\boldsymbol{\xi}$. That the spectral method can yield an estimate $\boldsymbol{x}_1$ with a positive ρ in high dimensional settings is a nontrivial property. To see this, assume that $\boldsymbol{\xi}$ is pointing towards the "north pole" in the unit $(n-1)$-sphere $\mathcal{S}^{n-1}$, as illustrated in Figure 1. If we choose $\boldsymbol{x}_1$ uniformly at random from $\mathcal{S}^{n-1}$, then with high probability, the resulting correlation $\sqrt{\rho(\boldsymbol{\xi}, \boldsymbol{x}_1)}$ will be of order $\mathcal{O}(1/\sqrt{n})$. In other words, for large n, most of the uniform measure on $\mathcal{S}^{n-1}$ is concentrated within a very thin band of width $\mathcal{O}(1/\sqrt{n})$ near the "equator" of the sphere (see Figure 1.)

Our analysis reveals a phase transition phenomenon that occurs at certain critical values of the sampling ratio. In particular, there exist a lower and an upper threshold, denoted by $\alpha_{c,\min}$ and $\alpha_{c,\max}$, respectively, that mark the transitions between two very different phases.

(a) An *uncorrelated phase* takes place when the sampling ratio $\alpha < \alpha_{c,\min}$. Within this phase, the limiting value $\rho(\alpha) = 0$, meaning that the estimate from the spectral method is asymptotically uncorrelated with the target vector $\boldsymbol{\xi}$. In this case, the spectral method is not effective, as its estimate $\boldsymbol{x}_1$ is no better than a random guess drawn uniformly from the hypersphere $\mathcal{S}^{n-1}$.

(b) A *correlated phase* takes place when $\alpha > \alpha_{c,\max}$, with $\alpha_{c,\max}$ being the upper threshold. Within this phase, the limiting value $\rho(\alpha) > 0$. Geometrically, the estimate $\boldsymbol{x}_1$ (or its negative version $-\boldsymbol{x}_1$) will be concentrated on the surface of a right-circular cone (see Figure 1) whose generating lines make an angle $\theta = \arccos\left(\sqrt{\rho(\alpha)}\right)$ to the target vector $\boldsymbol{\xi}$. Moreover, $\rho(\alpha)$ tends to 1 as $\alpha \to \infty$.

In many signal estimation models that we have studied so far, the two thresholds coincide, *i.e.* $\alpha_{c,\min} = \alpha_{c,\max}$, meaning that the phase transition happens at a single critical value of the sampling ratio. However, it is indeed possible that $\alpha_{c,\min} < \alpha_{c,\max}$, in which case a finite number of correlated and uncorrelated phases alternative when α varies within the interval $(\alpha_{c,\min}, \alpha_{c,\max})$. A concrete example demonstrating this more complicated situation can be found in Section IV-C.

The above phase transition phenomenon also has implications in terms of the computational complexity of the spectral method. In a correlated phase, there is a nonzero gap between the largest and the second largest eigenvalues of $\boldsymbol{D}_m$. As a result, the leading eigenvector $\boldsymbol{x}_1$ can be efficiently computed by using power iterations on $\boldsymbol{D}_m$. In contrast, within an uncorrelated phase, the gap of the eigenvalues converges to zero, making power iterations ineffective.

The rest of the paper is organized as follows. After precisely laying out the various technical assumptions, we present in Section II the main results of this work, stated as Theorem 1 and Proposition 1. Examples and numerical simulations are also provided there to demonstrate and verify these analytical results. In particular, as a worked example, we derive a universal closed-form expression for the limiting values $\rho(\alpha)$ for all acquisition models that generate one-bit $\{0, 1\}$ measurements. We prove Theorem 1 in Section III. Key to our proof is a deterministic, fixed-point characterization of the squared cosine similarity $\rho(\boldsymbol{\xi}, \boldsymbol{x}_1)$, which is valid for any finite dimension n and for any *deterministic* sensing vectors $\{\boldsymbol{a}_i\}$. When specialized to Gaussian measurements, this fixed-point characterization allows us to connect our problem to a generalized version of the spiked population model (see, *e.g.*, [27]–[29]) studied in random matrix theory. In Section IV, we look more closely at the phase transition phenomenon

predicted by our asymptotic results and prove Proposition 1. Section V concludes the paper with discussions on possible generalizations and improvements of our results as well as their connections to related work in the literature.

Notations: To study the high-dimensional limit of the spectral initialization method, we shall consider a sequence of problem instances, indexed by the ambient dimension n. For each n, we seek to estimate an underlying signal denoted by $\boldsymbol{\xi}_n \in \mathbb{R}^n$. Formally, we should use $D_{m(n)}$ to denote the data matrix, where $m(n)$ is the number of measurements as a function of the dimension n. However, to lighten the notation, we shall simply write $\boldsymbol{D}_m$, keeping the dependence of m on n implicit. $\boldsymbol{x}_1^n$ stands for a leading eigenvector of $\boldsymbol{D}_m$. We use $\xrightarrow{\mathcal{P}}$ and $\xrightarrow{\text{a.s.}}$ to denote convergence in probability and almost sure convergence, respectively. Let $\boldsymbol{M}$ be a symmetric matrix. Its eigenvalues in descending order are written as $\lambda_1^{\boldsymbol{M}} \geq \lambda_2^{\boldsymbol{M}} \geq \ldots \geq \lambda_n^{\boldsymbol{M}}$. In particular, $\lambda_1^{\boldsymbol{M}}$, sometimes also written as $\lambda_1(\boldsymbol{M})$, denotes the largest eigenvalue of $\boldsymbol{M}$. Throughout the paper, $\boldsymbol{A}^{1/2}$ stands for the principal square root of a positive semi-definite symmetric matrix $\boldsymbol{A}$. For any $a, b \in \mathbb{R}$, we write $\max\{a, b\}$ as $a \vee b$. Finally, $\mathbb{1}_{\mathcal{I}}(x)$ stands for the indicator function of a set $\mathcal{I}$.

II. Main Results

A. Technical Assumptions

In what follows, we first state the basic assumptions under which our results are proved.

(A.1) The sensing vectors are independent Gaussian random vectors. Specifically, let (a_{ij}), for $i, j \geq 1$, be a doubly infinite array of i.i.d. standard normal random variables. Then the ith sensing vector $\boldsymbol{a}_i = [a_{i1}, a_{i2}, \ldots, a_{in}]^T$.

(A.2) $m = m(n)$ with $\alpha_n = m(n)/n \to \alpha > 0$ as $n \to \infty$.

(A.3) $\|\boldsymbol{\xi}_n\| = \kappa > 0$.

(A.4) Let s, y and z be three random variables such that

$$s \sim \mathcal{N}(0,1),\ \mathbb{P}(y \,|\, s) = f(y \,|\, \kappa s),\ \text{and } z = \mathcal{T}(y), \quad (5)$$

where $f(\cdot\,|\,\cdot)$ is the conditional density function (1) associated with the observation model, and $\mathcal{T}(\cdot)$ is the preprocessing step used in the construction of $\boldsymbol{D}_m$ in (3). We shall assume that the probability measure of the random variable z is supported within a finite interval $[0, \tau]$. Throughout the paper, we always take τ to be the *tightest* such upper bound.

(A.5) As λ approaches τ from the right,

$$\lim_{\lambda \to \tau^+} \mathbb{E}\frac{z}{(\lambda - z)^2} = \lim_{\lambda \to \tau^+} \mathbb{E}\frac{zs^2}{\lambda - z} = \infty. \quad (6)$$

(A.6) The random variables z and s^2 are positively correlated: $\operatorname{cov}(z, s^2) = \mathbb{E}\, zs^2 - \mathbb{E}z\, \mathbb{E}s^2 > 0$, which is equivalent to

$$\mathbb{E}\, zs^2 > \mathbb{E}z. \quad (7)$$

The last three assumptions require some explanations. First, we note that assumption (A.4) is not particularly restrictive. The actual observations $\{y_i\}$ need not be bounded or nonnegative. The requirement that z should take values within a finite interval on the positive axis can be enforced by choosing a suitable function $\mathcal{T}(\cdot)$. For example, in the problem of phase retrieval, the measurement model ($y = s^2$) leads to unbounded $\{y_i\}$. We can choose

$$z = \mathcal{T}(y) = y\, \mathbb{1}_{[0,t]}(y), \quad (8)$$

where $t > 0$ is some parameter and $\mathbb{1}(\cdot)$ is the indicator function. This is indeed the trimming strategy proposed in [13]. As shown there, this boundedness condition on z is an essential ingredient in achieving linear sample complexities.

In assumption (A.5), the expressions in (6) essentially require that the random variable z should have sufficient probability mass near the upper bound τ. Let $h(z) = \mathbb{E}_{s|z}(s^2|z)$. We show in Appendix A that (6) holds when there exist some positive constants c_0 and ε such that the probability density function $p_Z(z)$ of z and the conditional moment $h(z)$ are both bounded below by c_0 for all $z \in [\tau - \varepsilon, \tau]$. The model in (8) represents one such case. Another sufficient condition for (6) to hold is when the law of z has a point mass at τ. The acquisition models described in (27) and (28) in later sections are examples for which this condition is applicable.

The inequality in (A.6) is also a natural requirement. To see this, we note that the data matrix $\boldsymbol{D}_m$ in (3) is the sample average of m i.i.d. random rank-one matrices $\left\{y_i \boldsymbol{a}_i \boldsymbol{a}_i^T\right\}_{i \leq m}$. When the number of samples m is large, this sample average should be "close" to the statistical expectation, *i.e.*,

$$\boldsymbol{D}_m \approx \mathbb{E}(z_i\, \boldsymbol{a}_i \boldsymbol{a}_i^T), \quad (9)$$

where $z_i \overset{\text{def}}{=} \mathcal{T}(y_i)$. To compute the above expectation, it will be convenient to assume that the underlying signal $\boldsymbol{\xi} = \kappa \boldsymbol{e}_1$, where $\boldsymbol{e}_1$ is the first vector of the canonical basis of $\mathbb{R}^n$. (This assumption can be made without loss of generality, due to the rotational invariance of the multivariate normal distribution.) Correspondingly, we can partition each sensing vector into two parts, as

$$\boldsymbol{a}_i^T = \begin{bmatrix} s_i & \boldsymbol{u}_i^T \end{bmatrix}, \quad (10)$$

so that $\boldsymbol{a}_i^T \boldsymbol{\xi} = \kappa s_i$ and the conditional density of y_i given s_i is $f(y \,|\, \kappa s_i)$. Since s_i, y_i and z_i are all independent of $\boldsymbol{u}_i$,

$$\begin{aligned} \mathbb{E}(z_i\, \boldsymbol{a}_i \boldsymbol{a}_i^T) &= \mathbb{E}\begin{bmatrix} z_i s_i^2 & z_i s_i \boldsymbol{u}_i^T \\ z_i s_i \boldsymbol{u}_i & z_i \boldsymbol{u}_i \boldsymbol{u}_i^T \end{bmatrix} \\ &= \begin{bmatrix} \mathbb{E}\, zs^2 & \boldsymbol{0} \\ \boldsymbol{0} & \mathbb{E}z\, \boldsymbol{I}_{n-1} \end{bmatrix}, \end{aligned} \quad (11)$$

where $\boldsymbol{I}_{n-1}$ is the identity matrix of size $(n-1)$. If the inequality $\mathbb{E}\, zs^2 > \mathbb{E}z$, as required in (A5), indeed holds, the leading eigenvector of the expected matrix will be $\boldsymbol{e}_1$, which is perfectly aligned with the target vector $\boldsymbol{\xi}$. Now since the data matrix $\boldsymbol{D}_m$ is an approximation of the expectation, the sample eigenvector should also be an approximation of $\boldsymbol{\xi}$.

The above argument provides an intuitive but nonrigorous explanation for why the spectral initialization method would work. The approximation in (9) can be made exact if the signal dimension n is kept fixed and the number of measurement m goes to infinity. However, we consider the case when m and n both tend to infinity, at a constant ratio $\alpha = m/n$ bounded away from 0 and ∞. In this regime, the approximation in (9) will not become an equality even if $m \to \infty$. As we will show,

the correlation $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ between the target vector $\boldsymbol{\xi}_n$ and the sample eigenvector $\boldsymbol{x}_1^n$ will converge to a function $\rho(\alpha)$.

A notable exception to (7) is when

$$g(s) \overset{\text{def}}{=} \mathbb{E}_{z|s}(z|s) \tag{12}$$

is an odd function plus some arbitrary constant C. In this case, $\mathbb{E}zs^2 = \mathbb{E}[g(s)s^2] = C$ and $\mathbb{E}z = \mathbb{E}g(s) = C$ and thus (7) does not hold. In practice, this means that the spectral method will not be effective for acquisition models such as $z = \text{sign}(s) + C$. We will revisit this point in Section V where we describe an alternative initialization scheme that can handle such cases.

As a final remark before we present our main results, we note that, since the eigenvector $\boldsymbol{x}_1^n$ is always normalized, the spectral method cannot provide any information about the norm of $\boldsymbol{\xi}_n$. However, in many cases when the sensing vectors are drawn from certain random ensembles, there are simple methods to accurately estimate $\kappa = \|\boldsymbol{\xi}_n\|$. We provide some discussions on how to do this in Appendix B.

B. Main Results: Asymptotic Characterizations

In this section, we summarize the main results of our work on an asymptotic characterization of the spectral method with Gaussian measurements. To state our results, we first need to introduce several helper functions. Let s, z be the random variables defined in (5). We consider two functions

$$\phi(\lambda) \overset{\text{def}}{=} \lambda\, \mathbb{E}\frac{zs^2}{\lambda - z} \tag{13}$$

and

$$\psi_\alpha(\lambda) \overset{\text{def}}{=} \lambda \left(1/\alpha + \mathbb{E}\frac{z}{\lambda - z}\right), \tag{14}$$

both defined on the open interval (τ, ∞), where τ is the bound in assumption (A.4). Within their domains, it is easy to check that both functions are convex. In particular, $\psi_\alpha(\lambda)$ achieves its minimum at a unique point denoted by

$$\overline{\lambda}_\alpha \overset{\text{def}}{=} \underset{\lambda > \tau}{\arg\min}\ \psi_\alpha(\lambda). \tag{15}$$

Finally, let

$$\zeta_\alpha(\lambda) \overset{\text{def}}{=} \psi_\alpha\big(\lambda \vee \overline{\lambda}_\alpha\big) \tag{16}$$

be a modification of $\psi_\alpha(\lambda)$. This new function is again defined for $\lambda \in (\tau, \infty)$.

Theorem 1: Under (A.1) – (A.6), the following hold:

1) There is a unique solution, denoted by λ_α^*, to the equation
$$\zeta_\alpha(\lambda) = \phi(\lambda), \qquad \lambda > \tau. \tag{17}$$
2) As $n \to \infty$,
$$\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n) \xrightarrow{\mathcal{P}} \begin{cases} 0, & \text{if } \psi_\alpha'(\lambda_\alpha^*) < 0, \\ \frac{\psi_\alpha'(\lambda_\alpha^*)}{\psi_\alpha'(\lambda_\alpha^*) - \phi'(\lambda_\alpha^*)}, & \text{if } \psi_\alpha'(\lambda_\alpha^*) > 0, \end{cases} \tag{18}$$
where $\psi_\alpha'(\cdot)$ and $\phi'(\cdot)$ denote the derivatives of the two functions.
3) Let $\lambda_1^{\boldsymbol{D}_m} \geq \lambda_2^{\boldsymbol{D}_m}$ be the top two eigenvalues of $\boldsymbol{D}_m$.
$$\lambda_1^{\boldsymbol{D}_m} \xrightarrow{\mathcal{P}} \zeta_\alpha(\lambda_\alpha^*) \quad \text{and} \quad \lambda_2^{\boldsymbol{D}_m} \xrightarrow{\mathcal{P}} \zeta_\alpha(\overline{\lambda}_\alpha) \tag{19}$$
as $n \to \infty$. Moreover, $\zeta_\alpha(\lambda_\alpha^*) \geq \zeta_\alpha(\overline{\lambda}_\alpha)$, with the inequality becoming strict *if and only if* $\psi_\alpha'(\lambda_\alpha^*) > 0$.

Remark 1: The above theorem, whose proof is given in Section III, provides a complete asymptotic characterization of the performance of the spectral method. In particular, the theorem shows that the squared cosine similarity $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ converges in probability to a deterministic value in the high-dimensional limit. Moreover, there exists a generic phase transition phenomenon: depending on the sign of the derivative $\psi_\alpha'(\cdot)$ at λ_α^*, the limiting value can be either zero (*i.e.*, the uncorrelated phase) or strictly positive (*i.e.*, the correlated phase.) The computational complexity of the spectral method is also very different in the two phases. Within the uncorrelated phase, the gap between the top two leading eigenvalues, $\lambda_1^{\boldsymbol{D}_m}$ and $\lambda_2^{\boldsymbol{D}_m}$, diminishes to zero, making iterative algorithms such as power iterations increasingly difficult to converge. In contrast, within the correlated phase, the spectral gap converges to a positive value.

It will be more convenient to characterize the above phase transitions in terms of the sampling ratio α. To do so, we first introduce a set Λ, containing all the zero-crossings of the function $\Delta(\lambda) = \lambda\, \mathbb{E}\frac{z}{(\lambda - z)^2} - \mathbb{E}\frac{zs^2}{\lambda - z}$ on the open interval (τ, ∞). We can show that Λ is always nonempty and that it contains a finite number of points (see Lemma 2 in Section IV-A.) Let

$$\lambda_{c,\min} \overset{\text{def}}{=} \min_{\lambda \in \Lambda} \lambda \quad \text{and} \quad \lambda_{c,\max} \overset{\text{def}}{=} \max_{\lambda \in \Lambda} \lambda$$

denote the smallest and the largest elements in Λ, respectively.

Proposition 1: Under (A.1) – (A.6), and as $n \to \infty$,

$$\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n) \xrightarrow{\mathcal{P}} \begin{cases} 0, & \text{if } \alpha < \alpha_{c,\min}, \\ \rho(\alpha), & \text{if } \alpha > \alpha_{c,\max}, \end{cases}$$

where

$$\alpha_{c,\min}^{-1} = \mathbb{E}\frac{z^2}{(\lambda_{c,\min} - z)^2},\ \alpha_{c,\max}^{-1} = \mathbb{E}\frac{z^2}{(\lambda_{c,\max} - z)^2}, \tag{20}$$

and $\rho(\alpha)$ is a function with the following parametric representation in terms of a parameter λ:

$$1/\alpha = \mathbb{E}\frac{zs^2 - z}{\lambda - z} \tag{21}$$

$$1/\rho = 1 + \left(\mathbb{E}\frac{zs^2 - z}{\lambda - z} - \mathbb{E}\frac{z^2}{(\lambda - z)^2}\right)^{-1} \mathbb{E}\frac{z^2 s^2}{(\lambda - z)^2}, \tag{22}$$

for all $\lambda > \lambda_{c,\max}$. Moreover, $\rho(\alpha) \to 1$ as $\alpha \to \infty$.

Remark 2: In many of the signal acquisition models we have studied, the set Λ contains exactly one element. In this case, $\lambda_{c,\min} = \lambda_{c,\max}$ and hence $\alpha_{c,\min} = \alpha_{c,\max}$. Consequently, the phase transition of the spectral method takes place at a single threshold value α_c, which separates the uncorrelated phase from the correlated one. However, it is indeed possible to find cases for which $\alpha_{c,\min} < \alpha_{c,\max}$. This leads to a more complicated scenario, where a finite number of correlated and uncorrelated phases can alternatively take place within the interval $(\alpha_{c,\min}, \alpha_{c,\max})$. One such example can be found in Section IV-C.

C. Worked-Example: Binary Models

To illustrate the results presented above, we consider here a special case where z_i takes only binary values $\{0,1\}$. This situation naturally appears in problems such as logistic regression and one-bit quantized sensing, where the measurements $y_i \in \{0,1\}$ and we can set $z_i = y_i$. For cases where the measurements $\{y_i\}$ are not necessarily binary, this type of one-bit model is still relevant whenever the preprocessing function $z = \mathcal{T}(x)$ generates binary outputs. The simplicity of this setting allows us to obtain closed-form expressions for the various quantities in Theorem 1 and Proposition 1.

To proceed, we first explicitly compute the functions $\phi(\lambda)$ and $\psi_\alpha(\lambda)$ defined in Section II-B as

$$\phi(\lambda) = \frac{c\lambda}{\lambda - 1} \quad \text{and} \quad \psi_\alpha(\lambda) = \lambda\Big(1/\alpha + \frac{d}{\lambda - 1}\Big),$$

where

$$c \overset{\text{def}}{=} \mathbb{E}\, zs^2 \quad \text{and} \quad d \overset{\text{def}}{=} \mathbb{E}\, z \tag{23}$$

and both functions are defined on the interval $\lambda > 1$. The minimum of $\psi_\alpha(\lambda)$ is achieved as $\overline{\lambda}_\alpha = 1 + \sqrt{\alpha d}$, and thus

$$\zeta_\alpha(\lambda) = \begin{cases} \lambda/\alpha + \lambda d/(\lambda - 1), & \text{for } \lambda \geq 1 + \sqrt{\alpha d} \\ (\sqrt{d} + 1/\sqrt{\alpha})^2, & \text{for } 1 < \lambda < 1 + \sqrt{\alpha d}. \end{cases}$$

Solving equation (17) and using (18), we get

$$\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n) \xrightarrow{\mathcal{P}} \begin{cases} 0, & \text{for } \alpha < \alpha_c, \\ \frac{\alpha - d/(c-d)^2}{\alpha + 1/(c-d)}, & \text{for } \alpha > \alpha_c, \end{cases} \tag{24}$$

where $\alpha_c = \frac{d}{(c-d)^2}$. (Note that this result can also be obtained by invoking the parametric characterization of $\rho(\alpha)$ given in Proposition 1.) Finally, the asymptotic predictions (19) for the top two eigenvalues can be computed as

$$\lambda_1^{\boldsymbol{D}_m} \xrightarrow{\mathcal{P}} \begin{cases} (\sqrt{d} + 1/\sqrt{\alpha})^2, & \text{for } \alpha < \alpha_c, \\ c + \frac{c}{\alpha(c-d)}, & \text{for } \alpha > \alpha_c, \end{cases} \tag{25}$$

and

$$\lambda_2^{\boldsymbol{D}_m} \xrightarrow{\mathcal{P}} (\sqrt{d} + 1/\sqrt{\alpha})^2 \tag{26}$$

for all α.

Remark 3: It is interesting to note that the asymptotic characterizations given in (24), (25) and (26) are *universal*, in the sense that they only depend on the two constants c and d defined in (23) but not on the exact details of the joint probability distributions of s, y and z. Thus, for one-bit models, it suffices to compute the constants in (23), which then completely determine the asymptotic performance of the spectral method.

D. Numerical Simulations

Example 1 (Logistic regression): Consider the case where $\{y_i\}$ are binary random variables generated according to the following conditional distribution:

$$f(y \,|\, \boldsymbol{a}^T\boldsymbol{\xi}_n) \sim \text{Bernoulli}\left(\frac{1}{1 + \exp\left\{-\boldsymbol{a}^T\boldsymbol{\xi}_n - \beta\right\}}\right), \tag{27}$$

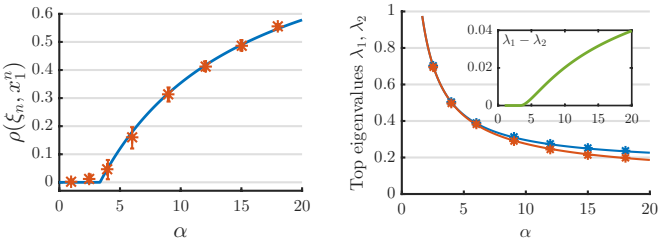

(a) Squared cosine similarity (b) Top two eigenvalues

Fig. 2. Analytical predictions v.s. numerical simulations for the binary logistic model in (27). Numerical results are averaged over 16 independent trials.

where β is some constant. Let $z_i = \mathcal{T}(y_i) = y_i$. Since $z_i \in \{0,1\}$, we just need to compute the constants c and d in (23), after which we can use the closed-form expressions (24), (25) and (26) to obtain the asymptotic predictions. In Figure 2(a) we compare the analytical prediction (24) of the squared cosine similarity with results of numerical simulations. In our experiment, we set the signal dimension to $n = 4096$. The norm of $\boldsymbol{\xi}_n$ is $\kappa = 3$, and $\beta = 6$. The sample averages and error bars (corresponding to one standard deviation) shown in the figure are calculated over 16 independent trials. We can see that the analytical predictions match numerical results very well. Figure 2(b) shows the top two eigenvalues. When $\alpha < \alpha_c$, the two eigenvalues are asymptotically equal, but they start to diverge as α becomes larger than α_c. To clearly illustrate this phenomenon, we plot in the insert the eigengap $\lambda_1 - \lambda_2$ as a function of α.

Example 2 (Phase retrieval): In the second example, we consider the problem of phase retrieval. For simplicity, we assume the noiseless setting, where $y_i = (\boldsymbol{a}_i^T\boldsymbol{\xi}_n)^2$. In [13], the authors show that it is important to omit large values of $\{y_i\}$, and they propose to use the scheme in (8) when constructing the data matrix $\boldsymbol{D}_m$. A different strategy can be found in [15], where the authors propose to use

$$z_i = \mathbb{1}(y_i > t). \tag{28}$$

In what follows, we shall refer to (8) and (28) as the trimming algorithm and the subset algorithm, respectively. Figure 3(a) shows the asymptotic performance of these two algorithms and compare them with numerical results ($n = 4096$ and 16 independent trials.) The performance of the subset algorithm (for which we choose the parameter $t = 1.5$) can be characterized by the closed-form formula (24). The trimming algorithm (for which we use $t = 3$) is more complicated as z_i is no longer binary. We use the parametric characterization in Proposition 1 to obtain its asymptotic performance. Again, our analytical predictions match numerical results. The performance of both algorithms clearly depends on the choice of the thresholding parameter t. To show this, we plot in Figure 3(b) the critical phase transition points α_c of both algorithms as functions of t. This points to the possibility of using our analytical prediction to optimally tune the algorithmic parameters. More generally, one can potentially use our analytical results to design the optimal preprocessing function $\mathcal{T}(\cdot)$. We leave this as an interesting avenue for future research.

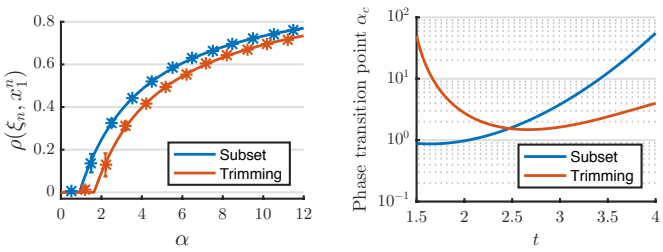

(a) Squared cosine similarity (b) Phase transition points

Fig. 3. (a) Analytical predictions v.s. numerical simulations for two different algorithms for phase retrieval. (b) The critical sampling ratio α_c of the two algorithms as functions of the threshold value t.

III. Proof of the Main Results

In this section, we prove Theorem 1, which provides an exact characterization of the asymptotic performance of the spectral method for signal estimation.

A. Overview

We first rewrite the data matrix $\boldsymbol{D}_m$ in (3) as

$$\boldsymbol{D}_m = \tfrac{1}{m}\boldsymbol{A}\boldsymbol{Z}\boldsymbol{A}^T, \tag{29}$$

where $\boldsymbol{A} = [\boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_m]$ is an $n \times m$ matrix of i.i.d. normal random variables and

$$\boldsymbol{Z} \stackrel{\text{def}}{=} \text{diag}\,\{z_1, z_2, \ldots, z_m\} \tag{30}$$

is a diagonal matrix with entries $z_i = \mathcal{T}(y_i)$. Our goal boils down to studying the largest eigenvalue of $\boldsymbol{D}_m$ and the associated eigenvector $\boldsymbol{x}_1^n$. To simplify notation, we shall first assume that $\boldsymbol{\xi}_n = \kappa \boldsymbol{e}_1$, with $\boldsymbol{e}_1$ being the first vector in the canonical basis.

Remark 4: The non-null eigenvalues of $\boldsymbol{D}_m$ are equal to those of a companion matrix

$$\widetilde{\boldsymbol{D}}_m = \tfrac{1}{m}\boldsymbol{Z}^{1/2}\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{Z}^{1/2},$$

which bears strong resemblance to a sample covariance matrix. Limiting spectral distributions (LSDs) of sample covariance matrices have been extensively studied in random matrix theory (see for instance [30] and the references given there.) As a special case, when $\boldsymbol{Z}$ is the identity matrix, the LSD of $\widetilde{\boldsymbol{D}}_m$ is given by the classical Marčenko-Pastur law [31]. Results for more general diagonal matrices $\boldsymbol{Z}$ are also available [32]. However, in these studies, $\boldsymbol{Z}$ and $\boldsymbol{A}$ need to be independent. A challenge in our problem is that $\boldsymbol{Z}$ and $\boldsymbol{A}$ are *correlated*. To see this, we partition each sensing vector $\boldsymbol{a}_i$ into two parts as in (10). We can then write

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{s}^T \\ \boldsymbol{U} \end{bmatrix}, \tag{31}$$

where $\boldsymbol{s} \stackrel{\text{def}}{=} [s_1, s_2, \ldots, s_m]^T$ is an m-dimensional Gaussian random vector, and $\boldsymbol{U}$ is an $(n-1) \times m$ matrix consisting of i.i.d. standard normal random variables. Since $\boldsymbol{\xi}_n = \kappa \boldsymbol{e}_1$, the diagonal elements of $\boldsymbol{Z}$ are independent of $\boldsymbol{U}$ but they do depend on $\boldsymbol{s}$ through $y_i \sim f(y \,|\, \kappa s_i)$. Consequently, we cannot apply existing results on the LSD of sample covariance matrices to our case.

Our proof of Theorem 1 consists of two main ingredients. First, we will show in Proposition 2 that $\lambda_1^{\boldsymbol{D}_m}$ and $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ can be obtained from a fixed-point equation involving a function $L_m(\mu)$, to be defined in (36), where $\mu > 0$ is an auxiliary variable. The main benefit of introducing the variable μ and the function $L_m(\mu)$ is that, for each $\mu > 0$, the above-mentioned correlation between $\boldsymbol{A}$ and $\boldsymbol{Z}$ can be effectively decoupled. This then allows us to obtain the second ingredient of our proof: using results from random matrix theory [29], [33], we show in Section III-C that $L_m(\mu)$, under the assumption of Gaussian sensing vectors, will converge almost surely to a deterministic limit function as the dimension $n \to \infty$ (see Proposition 4.)

B. A Fixed-Point Characterization

By substituting (31) into (29), we can write $\boldsymbol{D}_m$ in a more compact block-partitioned form as

$$\boldsymbol{D}_m = \left[\begin{array}{c|c} a_m & \boldsymbol{q}_m^T \\ \hline \boldsymbol{q}_m & \boldsymbol{P}_m \end{array}\right], \tag{32}$$

where

$$a_m \stackrel{\text{def}}{=} \tfrac{1}{m}\sum_{i=1}^m z_i s_i^2 \tag{33}$$

is a scalar that converges to $\mathbb{E}zs^2$ as $m \to \infty$,

$$\boldsymbol{P}_m \stackrel{\text{def}}{=} \tfrac{1}{m}\boldsymbol{U}\boldsymbol{Z}\boldsymbol{U}^T \tag{34}$$

is a symmetric matrix, and

$$\boldsymbol{q}_m \stackrel{\text{def}}{=} \tfrac{1}{m}\boldsymbol{U}\boldsymbol{v} \;\text{ with }\; \boldsymbol{v} \stackrel{\text{def}}{=} [z_1 s_1, z_2 s_2, \ldots, z_m s_m]^T. \tag{35}$$

Next, we consider a parametric family of matrices $\{\boldsymbol{P}_m + \mu \boldsymbol{q}_m \boldsymbol{q}_m^T : \mu > 0\}$, and let $L_m(\mu)$ denote their largest eigenvalues, *i.e.*,

$$L_m(\mu) \stackrel{\text{def}}{=} \lambda_1(\boldsymbol{P}_m + \mu \boldsymbol{q}_m \boldsymbol{q}_m^T). \tag{36}$$

In what follows, we show how to compute $\lambda_1^{\boldsymbol{D}_m}$ and $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ via a fixed-point equation involving $L_m(\mu)$. Since we assume that $\boldsymbol{\xi} = \kappa\, \boldsymbol{e}_1$ and that the leading eigenvector $\boldsymbol{x}_1^n$ is normalized, the quantity $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ is equal to $(\boldsymbol{e}_1^T \boldsymbol{x}_1^n)^2$, the squared magnitude of the first element of the eigenvector.

Our discussions below are general and they apply to any block-partitioned matrix in the form

$$\boldsymbol{D} = \left[\begin{array}{c|c} a & \boldsymbol{q}^T \\ \hline \boldsymbol{q} & \boldsymbol{P} \end{array}\right]. \tag{37}$$

Its components $a \in \mathbb{R}$, $\boldsymbol{P} \in \mathbb{R}^{(n-1)\times(n-1)}$ and $\boldsymbol{q} \in \mathbb{R}^{n-1}$ can be arbitrarily chosen, not necessarily defined as in (33), (34) and (35). Our only requirements are that $\boldsymbol{P}$ is a symmetric matrix and that $\|\boldsymbol{q}\| \neq 0$.

Let $\lambda_1^{\boldsymbol{P}} \geq \lambda_2^{\boldsymbol{P}} \geq \ldots \lambda_{n-1}^{\boldsymbol{P}}$ be the set of eigenvalues of $\boldsymbol{P}$, and let $\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{n-1}$ be a corresponding set of orthonormal eigenvectors. Consider a function

$$R(\lambda) \stackrel{\text{def}}{=} \boldsymbol{q}^T(\boldsymbol{P} - \lambda \boldsymbol{I})^{-1}\boldsymbol{q} = \sum_{i=1}^{n-1} \frac{(\boldsymbol{w}_i^T \boldsymbol{q})^2}{\lambda_i^{\boldsymbol{P}} - \lambda}, \tag{38}$$

which has poles on those eigenvalues for which $\boldsymbol{w}_i^T\boldsymbol{q} \neq 0$. In what follows, we restrict the domain of $R(\lambda)$ to

$$\lambda \in \left(\max\left\{\lambda_i^{\boldsymbol{P}} : \boldsymbol{w}_i^T\boldsymbol{q} \neq 0\right\}, \infty\right).$$

Within this open interval, $R(\lambda)$ is a well-defined smooth function. It increases monotonically from $-\infty$ to 0, and thus it admits a functional inverse, denoted by $R^{-1}(x)$, for all $x < 0$. Similar to (36), we define

$$L(\mu) = \lambda_1(\boldsymbol{P} + \mu\boldsymbol{q}\boldsymbol{q}^T)$$

for all $\mu > 0$.

Lemma 1: Let $\boldsymbol{P}$ be a symmetric matrix and $\boldsymbol{q}$ a nonzero vector. Then, for each $\mu > 0$,

$$L(\mu) = R^{-1}(-1/\mu) \vee \lambda_1^{\boldsymbol{P}}. \tag{39}$$

Moreover, $L(\mu)$ is a nondecreasing convex function with $\lim_{\mu\to\infty} L(\mu) = \infty$. It is differentiable everywhere on $(0, \infty)$ except at (up to) one point.

Proof: Since $\boldsymbol{P}$ is diagonalizable by an orthonormal matrix, we can assume without loss of generality that $\boldsymbol{P}$ is a diagonal matrix. In this case, we can simply write $R(\lambda) = \sum_i \frac{q_i^2}{\lambda_i^{\boldsymbol{P}} - \lambda}$, and this function is defined on the open interval $(\max\left\{\lambda_i^{\boldsymbol{P}} : q_i \neq 0\right\}, \infty)$.

Using the matrix determinant lemma [34], we can compute the characteristic polynomial of $\boldsymbol{P} + \mu\boldsymbol{q}\boldsymbol{q}^T$ as

$$\begin{aligned} c(\lambda) &= \det(\lambda\boldsymbol{I} - \boldsymbol{P} - \mu\boldsymbol{q}\boldsymbol{q}^T) \\ &= \det(\lambda\boldsymbol{I} - \boldsymbol{P}) - \mu\boldsymbol{q}^T \operatorname{adj}(\lambda\boldsymbol{I} - \boldsymbol{P})\boldsymbol{q} \qquad (40) \\ &= \prod_i (\lambda - \lambda_i^{\boldsymbol{P}}) - \mu\sum_i q_i^2 \prod_{j\neq i}(\lambda - \lambda_j^{\boldsymbol{P}}). \qquad (41) \end{aligned}$$

In (40), $\operatorname{adj}(\cdot)$ stands for the adjugate of a matrix. To reach (41), we have used the fact that, for any diagonal matrix $\boldsymbol{A} = \operatorname{diag}\{d_1, d_2, \ldots, d_{n-1}\}$, $\operatorname{adj}(\boldsymbol{A}) = \operatorname{diag}\left\{\prod_{j\neq 1} d_j, \prod_{j\neq 2} d_j, \ldots, \prod_{j\neq n-1} d_j\right\}$.

Partition the set $\{1, 2, \ldots, n-1\}$ into two subsets:

$$\mathcal{I}_1 = \left\{\lambda_i^{\boldsymbol{P}} : q_i = 0\right\} \quad \text{and} \quad \mathcal{I}_2 = \left\{\lambda_i^{\boldsymbol{P}} : q_i \neq 0\right\}. \tag{42}$$

We observe that the characteristic polynomial can be factored into $c(\lambda) = c_1(\lambda)c_2(\lambda)$, where $c_1(\lambda) = \prod_{i\in\mathcal{I}_1}(\lambda_i^{\boldsymbol{P}} - \lambda)$ and

$$c_2(\lambda) = \prod_{i\in\mathcal{I}_2}(\lambda - \lambda_i^{\boldsymbol{P}}) - \mu\sum_{i\in\mathcal{I}_2} q_i^2 \prod_{j\in\mathcal{I}_2\setminus i}(\lambda - \lambda_j^{\boldsymbol{P}}).$$

It is possible that the first subset $\mathcal{I}_1$ is empty, in which case $c_1(\lambda)$ is understood to be equal to 1, but $\mathcal{I}_2$ is never empty, since $\boldsymbol{q} \neq \boldsymbol{0}$. Next, we study the largest root of the polynomial $c_2(\lambda)$. For any $\lambda > \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$, we can write

$$\begin{aligned} c_2(\lambda) &= \Big(1 - \mu\sum_{i\in\mathcal{I}_2} \frac{q_i^2}{\lambda - \lambda_i^{\boldsymbol{P}}}\Big) \prod_{i\in\mathcal{I}_2}(\lambda - \lambda_i^{\boldsymbol{P}}) \\ &= \big(1 + \mu R(\lambda)\big) \prod_{i\in\mathcal{I}_2}(\lambda - \lambda_i^{\boldsymbol{P}}). \qquad (43) \end{aligned}$$

Recall that $R(\lambda)$ is the function defined in (38) and $R^{-1}(\cdot)$ is its functional inverse. It follows from (43) that $R^{-1}(-1/\mu)$ is the only root of $c_2(\lambda)$ in the interval $(\max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}, \infty)$, and therefore it is also the largest root. Due to the factorization $c(\lambda) = c_1(\lambda)c_2(\lambda)$, we have

$$L(\mu) = \max\{\lambda_i : i \in \mathcal{I}_1\} \vee R^{-1}(-1/\mu).$$

Finally, since $R^{-1}(-1/\mu) > \max\{\lambda_i : i \in \mathcal{I}_2\}$, we reach the formula in (39).

By construction, $R^{-1}(-1/\mu)$ is strictly increasing and $\lim_{\mu\to\infty} R^{-1}(-1/\mu) = \infty$. It is also differentiable everywhere on $(0, \infty)$. It follows that $L(\mu)$ is nondecreasing with $L(\infty) = \infty$, and that the function is differentiable everywhere except for at most one point μ_0, which, if it exists, must satisfy the identity $R^{-1}(-1/\mu_0) = \lambda_1^{\boldsymbol{P}}$. Finally, the convexity of $L(\mu)$ follows from the fact that it is the maximum of a set of linear functions, as $L(\mu) = \lambda_1(\boldsymbol{P} + \mu\boldsymbol{q}\boldsymbol{q}^T) = \max_{\boldsymbol{x}:\|\boldsymbol{x}\|=1} \boldsymbol{x}^T(\boldsymbol{P} + \mu\boldsymbol{q}\boldsymbol{q}^T)\boldsymbol{x}$. ∎

Given a block-partitioned matrix $\boldsymbol{D}$, the following proposition shows that its leading eigenvalue $\lambda_1^{\boldsymbol{D}}$ and the squared cosine similarity $(\boldsymbol{e}_1^T\boldsymbol{x}_1)^2$ can be obtained from the function $L(\mu)$.

Proposition 2: Let $\mu^* > 0$ be the unique solution to the fixed-point equation

$$\mu = (L(\mu) - a)^{-1}. \tag{44}$$

Then, $\lambda_1^{\boldsymbol{D}} = L(\mu^*)$ and

$$(\boldsymbol{e}_1^T\boldsymbol{x}_1)^2 \in \left[\frac{\partial_- L(\mu^*)}{\partial_- L(\mu^*) + (1/\mu^*)^2}, \frac{\partial_+ L(\mu^*)}{\partial_+ L(\mu^*) + (1/\mu^*)^2}\right], \tag{45}$$

where $\partial_- L(\mu)$ and $\partial_+ L(\mu)$ denote the left and right derivatives of $L(\mu)$, respectively. In particular, if $L(\mu)$ is differentiable at μ^*, then

$$(\boldsymbol{e}_1^T\boldsymbol{x}_1)^2 = \frac{L'(\mu^*)}{L'(\mu^*) + (1/\mu^*)^2}. \tag{46}$$

Remark 5: We prove this result in Appendix C. Note that (44) is equivalent to

$$L(\mu) = a + 1/\mu. \tag{47}$$

Since $L(\mu)$ is nondecreasing with $L(\infty) = \infty$ whereas $a + 1/\mu$ decreases monotonically from ∞ to 0, the equation (47), and thus (44), always admits one and only one solution. Moreover, by Lemma 1, $L(\mu)$ is a convex function, and therefore its left and right derivatives always exist.

C. Asymptotic Limit of $L_m(\mu)$

The characterization given in Proposition 2 is valid for any block-partitioned matrix in the form of (37). When applied to the specific case of our data matrix in (32), with its components a_m, $\boldsymbol{P}_m$ and $\boldsymbol{q}_m$ defined as in (33), (34) and (35), this result provides a very general *deterministic* characterization of the performance of the spectral method that is valid for any finite dimension n and for any sensing vectors.

Next, we specialize to the case of i.i.d. Gaussian sensing vectors and show that $L_m(\mu)$ converges almost surely to a deterministic function as $m, n \to \infty$. To that end, we note that $L_m(\mu)$ is the leading eigenvalue of

$$\boldsymbol{P}_m + \mu\boldsymbol{q}_m\boldsymbol{q}_m^T = \tfrac{1}{m}\boldsymbol{U}\boldsymbol{M}_m\boldsymbol{U}^T, \tag{48}$$

where

$$\boldsymbol{M}_m \stackrel{\text{def}}{=} \boldsymbol{Z} + \tfrac{\mu}{m}\boldsymbol{v}\boldsymbol{v}^T \tag{49}$$

is a rank-one perturbation of the diagonal matrix $\boldsymbol{Z}$ given in (30). Since $\boldsymbol{U}$ and $\boldsymbol{M}_m$ are independent, we first study the spectrum of $\boldsymbol{M}_m$.

Let $\lambda_1^{\boldsymbol{M}_m} \geq \lambda_2^{\boldsymbol{M}_m} \geq \ldots \geq \lambda_m^{\boldsymbol{M}_m}$ be the set of eigenvalues of $\boldsymbol{M}_m$ in descending order. Let

$$f^{\boldsymbol{M}_m}(\lambda) \stackrel{\text{def}}{=} \frac{1}{m-1}\sum_{i=2}^{m}\delta(\lambda - \lambda_i^{\boldsymbol{M}_m})$$

be the empirical spectral measure of the last $m-1$ eigenvalues.

Proposition 3: Fix $\mu > 0$. As $m, n \to \infty$, the empirical spectral measure $f^{\boldsymbol{M}_m}(\lambda)$ converges almost surely to the probability law of the random variable z. Meanwhile,

$$\lambda_1^{\boldsymbol{M}_m} \xrightarrow{\text{a.s.}} Q^{-1}(1/\mu), \tag{50}$$

where $Q^{-1}(\cdot)$ is the functional inverse of the function

$$Q(\lambda) = \mathbb{E}\frac{z^2 s^2}{\lambda - z}. \tag{51}$$

The domain of $Q(\lambda)$ is the open interval (τ, ∞), with τ being the upper bound of the support of the probability law of z.

Remark 6: By construction, $Q(\lambda)$ is a continuous and strictly decreasing function with $Q(\infty) = 0$. Assumption (A.5) further guarantees that $\lim_{\lambda\to\tau^+} Q(\lambda) = \infty$. Thus, $Q(\lambda)$ admits a functional inverse and that $Q^{-1}(1/\mu)$ is well-defined for all $\mu > 0$.

According to assumption (A.4) stated in Section II-A, the law of z is supported within the interval $[0, \tau]$. The above proposition, whose proof can be found in Appendix D, shows that the spectrum of $\boldsymbol{M}_m$ consists of two parts: a "bulk spectrum" of $m-1$ eigenvalues supported within $[0, \tau]$ and a single spiked eigenvalue $\lambda_1^{\boldsymbol{M}_m}$ well separated from the bulk. This setting is a generalization of the classical spiked population model [27]. Adapting the results given in [29] (see also [33] for related results under more general settings), we thus reach the second important ingredient of our proof of Theorem 1, characterizing the asymptotic limit of $L_m(\mu)$.

Proposition 4: For each fixed $\mu > 0$,

$$L_m(\mu) \xrightarrow{\text{a.s.}} \zeta_\alpha(Q^{-1}(1/\mu)), \tag{52}$$

where $\zeta_\alpha(\cdot)$ is the function defined in (16) and $Q^{-1}(1/\mu)$ is the limit value in (50).

Proof: Recall from (48) that $L_m(\mu)$ is the leading eigenvalue of $\frac{1}{m}\boldsymbol{U}\boldsymbol{M}_m\boldsymbol{U}^T$. Since $\boldsymbol{U}$ and $\boldsymbol{M}_m$ are independent, and since $\boldsymbol{U}$ is a Gaussian random matrix with a rotationally invariant distribution, we can equivalently study the leading eigenvalue of the following matrix

$$\tfrac{1}{m}\boldsymbol{U}\,\mathrm{diag}\left\{\lambda_1^{\boldsymbol{M}_m}, \lambda_2^{\boldsymbol{M}_m}, \ldots, \lambda_m^{\boldsymbol{M}_m}\right\}\boldsymbol{U}^T. \tag{53}$$

Proposition 3 shows that $\left\{\lambda_i^{\boldsymbol{M}} : i \geq 2\right\}$ form a bulk spectrum, which converges to the law of z as $m \to \infty$, whereas $\lambda_1^{\boldsymbol{M}}$ converges to a "spike" $\lambda_\mu = Q^{-1}(1/\mu) > \tau$, which is separated from the bulk.

The asymptotic limits of extreme sample eigenvalues of matrices in the form of (53) have been studied in [29], [33]. In our proof, we use the asymptotic characterization given in [29]. Key to this asymptotic analysis is the function $\psi_\alpha(\lambda)$ defined[1] in (14). The asymptotic behaviors of the leading sample eigenvalue turn out to depend on the sign of $\psi'_\alpha(\lambda)$ at the point λ_μ:

In particular, applying [29, Theorem 4.1], we have

$$L_m(\mu) \xrightarrow{\text{a.s.}} \psi_\alpha(\lambda_\mu) \quad \text{if } \psi'_\alpha(\lambda_\mu) > 0. \tag{54}$$

The case when $\psi'_\alpha(\lambda_\mu) \leq 0$ is covered in [29, Theorem 4.2]. Adapting that result to our specific setting, we have

$$L_m(\mu) \xrightarrow{\text{a.s.}} \min_{\lambda>\tau}\psi_\alpha(\lambda) \quad \text{if } \psi'_\alpha(\lambda_\mu) \leq 0. \tag{55}$$

As an equivalent form, we can write $\psi_\alpha(\lambda) = \mathbb{E}z + \frac{\lambda}{\alpha} + \mathbb{E}\frac{z^2}{\lambda - z}$. From this, we can easily check that $\psi_\alpha(\lambda)$ is a convex function and that it admits a unique minimum within its domain (τ, ∞). It follows that the two separate cases in (54) and (55) can be more compactly written as $L_m(\mu) \xrightarrow{\text{a.s.}} \zeta_\alpha(\lambda_\mu)$, where $\zeta_\alpha(\cdot)$ is the modified function defined in (16). ∎

D. Proof of Theorem 1

We are now ready to prove our asymptotic characterizations given in Theorem 1. Since the sensing vectors $\boldsymbol{a}_i$ are drawn from the rotationally invariant multivariate normal distribution, the quantity $\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n)$ for a general vector $\boldsymbol{\xi}_n$ (with $\|\boldsymbol{\xi}_n\| = \kappa$) and $\rho(\kappa\boldsymbol{e}_1, \boldsymbol{x}_1^n)$ for the special case $\boldsymbol{\xi}_n = \kappa\boldsymbol{e}_1$ have exactly the same probability distribution. In what follows, we will carry out the proof by assuming that the target vector $\boldsymbol{\xi}_n = \kappa\boldsymbol{e}_1$. By showing that $\rho(\kappa\boldsymbol{e}_1, \boldsymbol{x}_1^n)$ converges to the right-hand side of (18) *almost surely*, the convergence to the same limit in probability for a general $\boldsymbol{\xi}_n$ then follows as an immediate consequence.

To start, we use the deterministic characterization given in Proposition 2. For each $m \geq 1$, let μ_m be the unique fixed-point of (44). Equivalently, μ_m satisfies the identity

$$L_m(\mu_m) - 1/\mu_m = a_m.$$

By Proposition 3, for every fixed μ,

$$L_m(\mu) - 1/\mu \xrightarrow{\text{a.s.}} \zeta_\alpha(Q^{-1}(1/\mu)) - 1/\mu \tag{56}$$

as $m \to \infty$. Since $L_m(\mu)$ and $\zeta_\alpha(\mu)$ are nondecreasing, the two functions on both sides of (56) are strictly increasing. This condition, together with the fact that $a_m \xrightarrow{\text{a.s.}} \mathbb{E}zs^2$, allows us to apply Lemma 3 in Appendix E to conclude $\mu_m \xrightarrow{\text{a.s.}} \mu^*$, where μ^* is the unique point such that

$$\zeta_\alpha(Q^{-1}(1/\mu^*)) = \mathbb{E}zs^2 + 1/\mu^*. \tag{57}$$

To determine the asymptotic behavior of the leading eigenvector $\boldsymbol{x}_1^n$, we use the characterization given in (45). Since $\left\{L_m(\mu)\right\}$ are convex functions, we apply Lemma 4 in Appendix E. In particular, if $\zeta_\alpha(Q^{-1}(1/\mu))$ is differentiable at $\mu = \mu^*$, that lemma gives us

$$\partial_- L_m(\mu_m) \xrightarrow{\text{a.s.}} \left.\frac{\mathrm{d}\zeta_\alpha(Q^{-1}(1/\mu))}{\mathrm{d}\mu}\right|_{\mu^*} = \frac{-\zeta'_\alpha(Q^{-1}(1/\mu^*))}{Q'(Q^{-1}(1/\mu^*))(\mu^*)^2}$$

[1] We have adapted the original definition of $\psi_\alpha(\lambda)$ in [29, eq. (3.2)] because our matrix in (53) has a slightly different scaling from the one considered in [29].

and similarly

$$\partial_+ L_m(\mu_m) \xrightarrow{\text{a.s.}} \frac{-\zeta_\alpha'(Q^{-1}(1/\mu^*))}{Q'(Q^{-1}(1/\mu^*))(\mu^*)^2}.$$

Substituting these limits into (45), we get

$$(\boldsymbol{e}_1^T \boldsymbol{x}_1^n)^2 \xrightarrow{\text{a.s.}} \frac{\zeta_\alpha'(Q^{-1}(1/\mu^*))}{\zeta_\alpha'(Q^{-1}(1/\mu^*)) - Q'(Q^{-1}(1/\mu^*))}. \quad (58)$$

To simplify the above expression, we introduce a change of variable, writing $\lambda = Q^{-1}(1/\mu)$. In particular, $\lambda^* = Q^{-1}(1/\mu^*)$. Using the characterization (57) and recalling the definition of $Q(\lambda)$ in (51), we get

$$\zeta_\alpha(\lambda^*) = \mathbb{E} z s^2 + Q(\lambda^*) = \phi(\lambda^*), \quad (59)$$

where $\phi(\cdot)$ is defined in (13). By their constructions, it is easily checked that $\zeta_\alpha(\lambda)$ is a nondecreasing continuous function on (τ, ∞) whereas $\phi(\lambda)$ is a strictly decreasing continuous function. Moreover, by assumption (A.5), $\lim_{\lambda \to \tau^+} \phi(\lambda) = \infty$. Thus, the existence of λ^* satisfying (59) and its uniqueness are guaranteed. Substituting $\lambda^* = Q^{-1}(1/\mu^*)$ into (58) gives us

$$(\boldsymbol{e}_1^T \boldsymbol{x}_1^n)^2 \xrightarrow{\text{a.s.}} \frac{\zeta_\alpha'(\lambda^*)}{\zeta_\alpha'(\lambda^*) - \phi'(\lambda^*)},$$

where we have also used the fact that $Q'(\lambda) = \phi'(\lambda)$. To reach the characterization (18) given in the theorem, we just need to note that, by its definition in (16), $\zeta_\alpha'(\lambda) = \psi_\alpha'(\lambda)$ if $\psi_\alpha'(\lambda) > 0$ and $\zeta_\alpha'(\lambda) = 0$ if $\psi_\alpha'(\lambda) < 0$.

Next, we characterize the first two eigenvalues $\lambda_1^{\boldsymbol{D}_m}$ and $\lambda_2^{\boldsymbol{D}_m}$. By Proposition 2, the leading eigenvalue $\lambda_1^{\boldsymbol{D}_m} = L_m(\mu_m)$. Since $\mu_m \xrightarrow{\text{a.s.}} \mu^*$, applying Lemma 3 stated in Appendix E leads to

$$\lambda_1^{\boldsymbol{D}_m} \xrightarrow{\text{a.s.}} \zeta_\alpha(Q^{-1}(1/\mu^*)) = \zeta_\alpha(\lambda^*).$$

Recall from (32) that $\boldsymbol{P}_m$ is a principal submatrix of $\boldsymbol{D}_m$ obtained by deleting the first row and column of $\boldsymbol{D}_m$. It follows from the standard Cauchy interlacing theorem (see, *e.g.*, [35, Theorem 4.3.8]) that

$$\lambda_2^{\boldsymbol{P}_m} \le \lambda_2^{\boldsymbol{D}_m} \le \lambda_1^{\boldsymbol{P}_m} \quad (60)$$

Applying [29, Lemma 3.1] (which is due to [36]), the upper edge of the support of the limiting spectral density of $\boldsymbol{P}_m$ is given by

$$\min_{\lambda > \tau} \psi_\alpha(\lambda) = \zeta_\alpha(\overline{\lambda}_\alpha),$$

where $\overline{\lambda}_\alpha$ is the minimizing point defined in (15). It follows that $\lambda_2^{\boldsymbol{P}_m} \xrightarrow{\text{a.s.}} \zeta_\alpha(\overline{\lambda}_\alpha)$ and $\lambda_1^{\boldsymbol{P}_m} \xrightarrow{\text{a.s.}} \zeta_\alpha(\overline{\lambda}_\alpha)$, and thus

$$\lambda_2^{\boldsymbol{D}_m} \xrightarrow{\text{a.s.}} \zeta_\alpha(\overline{\lambda}_\alpha)$$

by the interlacing inequalities in (60). Finally, by the constructions of $\psi_\alpha(\lambda)$ and $\zeta_\alpha(\lambda)$, we have $\zeta_\alpha(\lambda) > \zeta_\alpha(\overline{\lambda}_\alpha)$ if and only if $\psi_\alpha'(\lambda) > 0$, and the proof is complete.

IV. Sampling Ratios and Phase Transitions

In this section, we study the phase transition phenomena characterized in Theorem 1 in more detail. In particular, we prove Proposition 1 (as stated in Section II-B), which specifies the phase transitions and the asymptotic limits of the cosine similarities in terms of the sampling ratio α.

A. Critical Sampling Ratios

By Theorem 1, whether the leading eigenvector $\boldsymbol{x}_1^n$ is asymptotically correlated or uncorrelated with the target vector $\boldsymbol{\xi}_n$ depends on the sign of the derivative $\psi_\alpha'(\lambda)$ evaluated at a point λ_α^*. And this point is uniquely defined through the equation $\zeta_\alpha(\lambda_\alpha^*) = \phi(\lambda_\alpha^*)$. Let $\overline{\lambda}_\alpha$, defined in (15), be the point at which the strictly convex function $\psi_\alpha(\lambda)$ achieves its minimum. Calculating the derivative of $\psi_\alpha(\lambda)$ and setting it to zero, we get

$$1/\alpha = \mathbb{E} \frac{z^2}{(\overline{\lambda}_\alpha - z)^2}. \quad (61)$$

By the construction of the function $\zeta_\alpha(\lambda)$ in (16) and by the monotonicity of $\phi(\lambda)$, we can conclude that $\psi'(\lambda_\alpha^*) > 0$ *if and only if*

$$\psi_\alpha(\overline{\lambda}_\alpha) < \phi(\overline{\lambda}_\alpha).$$

Substituting (61) into (14) gives us $\psi_\alpha(\overline{\lambda}_\alpha) = \overline{\lambda}_\alpha^2 \, \mathbb{E} \frac{z}{(\overline{\lambda}_\alpha - z)^2}$. Thus, transitions between the correlated and uncorrelated phases take place exactly at the zero-crossings of the function

$$\Delta(\lambda) = \mathbb{E} \frac{\lambda z}{(\lambda - z)^2} - \mathbb{E} \frac{z s^2}{\lambda - z}, \quad (62)$$

where $\Delta(\lambda)$ is obtained by removing a common factor $\overline{\lambda}_\alpha$ from the difference $\psi_\alpha(\overline{\lambda}_\alpha) - \phi(\overline{\lambda}_\alpha)$ and by writing $\overline{\lambda}_\alpha$ simply as λ. Let Λ be the set consisting of all the zero-crossings of $\Delta(\lambda)$ within the open interval (τ, ∞). Using (61), we can then establish a one-to-one mapping between points in Λ and a set of critical values of the sampling ratios.

Lemma 2: The set Λ is nonempty. It contains a finite number of points, denoted by $\lambda_{c,1} \le \lambda_{c,2} \le \ldots \le \lambda_{c,r}$ for some $r \ge 1$. Moreover,

$$\lambda_{c,r} \le \frac{\tau}{1 - \sqrt{\mathbb{E} z / \mathbb{E} z s^2}}. \quad (63)$$

Proof: We first show that Λ is nonempty. For $\lambda > \tau$, applying the Cauchy-Schwartz inequality gives us

$$\begin{aligned} \Delta(\lambda) &\ge \mathbb{E} \frac{\lambda z}{(\lambda - z)^2} - \left(\mathbb{E} \frac{z^2}{(\lambda - z)^2} \right)^{1/2} (\mathbb{E}\, s^4)^{1/2} \\ &\ge \mathbb{E} \frac{\tau z}{(\lambda - z)^2} - \left(\mathbb{E} \frac{3 \tau z}{(\lambda - z)^2} \right)^{1/2}. \end{aligned}$$

By assumption (A.5), $\mathbb{E} \frac{z}{(\lambda - z)^2} \to \infty$ as λ approaches τ from the right. Thus, we have

$$\lim_{\lambda \to \tau^+} \Delta(\lambda) = \infty. \quad (64)$$

To study the function $\Delta(\lambda)$ as $\lambda \to \infty$, we note that

$$\Delta(\lambda) \le \frac{1}{\lambda} \left(\frac{\lambda^2}{(\lambda - \tau)^2} \mathbb{E} z - \mathbb{E} z s^2 \right). \quad (65)$$

By assumption (A.6), $\mathbb{E} z < \mathbb{E} z s^2$. We can then conclude from inequality (65) that

$$\Delta(\lambda) < 0, \quad \text{for all sufficiently large } \lambda. \quad (66)$$

Since $\Delta(\lambda)$ is a continuous function, (64) and (65) imply that there must exist at least one zero-crossing.

Next, we show the upper bound given in (63). For any $\lambda_c \in \Lambda$, we have from (62) that

$$\lambda_c = \Big(\mathbb{E}\frac{zs^2}{\lambda - z}\Big)\Big(\mathbb{E}\frac{z}{(\lambda - z)^2}\Big)^{-1}. \tag{67}$$

By assumption (A.4), z is bounded within $[0, \tau]$. It follows that

$$\mathbb{E}\frac{zs^2}{\lambda - z} \geq \frac{\mathbb{E}zs^2}{\lambda} \quad \text{and} \quad \mathbb{E}\frac{z}{(\lambda - z)^2} \leq \frac{\mathbb{E}z}{(\lambda - \tau)^2}.$$

Substituting the above inequalities into (67) gives us $(\mathbb{E}z)\lambda_c^2 \geq (\mathbb{E}zs^2)(\lambda_c - \tau)^2$, which, after some simple manipulations, leads to the upper bound given in (63).

Finally, to show that Λ is a finite set, we extend $\Delta(\lambda)$ in (62) to the complex domain $\big\{\lambda \in \mathbb{C} : \mathrm{Re}(\lambda) > \tau\big\}$. Since $\Delta(\lambda)$ is analytic and it is not zero everywhere, by the principle of permanence, it has at most a finite number of zeros in the bounded domain $(\tau, \frac{\tau}{1-\sqrt{\mathbb{E}z/\mathbb{E}zs^2}})$. ∎

B. Proof of Proposition 1

Write $\lambda_{c,\min} = \lambda_{c,1}$ and $\lambda_{c,\max} = \lambda_{c,r}$. The corresponding critical sampling ratios $\alpha_{c,\min}$ and $\alpha_{c,\max}$, as defined in (20), are obtained through the one-to-one mapping given in (61).

Fix $\alpha < \alpha_{c,\min}$. By the monotonicity of the mapping (61), the corresponding $\overline{\lambda}_\alpha$ is strictly less than the smallest zero-crossing point $\lambda_{c,1}$. From the proof of Lemma 2, we conclude that $\Delta(\overline{\lambda}_\alpha) > 0$, and thus $\zeta_\alpha(\cdot)$ and $\phi_\alpha(\cdot)$ intersects at a point $\lambda_\alpha^* < \overline{\lambda}_\alpha$. This implies that $\psi'_\alpha(\lambda_\alpha^*) < 0$ and thus Theorem 1 gives us

$$\rho(\boldsymbol{\xi}_n, \boldsymbol{x}_1^n) \xrightarrow{\mathcal{P}} 0.$$

Now fix $\alpha > \alpha_{c,\max}$, in which case $\overline{\lambda}_\alpha > \lambda_{c,\max}$. Since $\Delta(\overline{\lambda}_\alpha) < 0$, we must have $\lambda_\alpha^* > \overline{\lambda}_\alpha$ and thus $\psi'_\alpha(\lambda_\alpha^*) > 0$. To derive the parametric form of $\rho(\alpha)$ given in the statement of the proposition, we note that

$$\zeta_\alpha(\lambda) = \psi_\alpha(\lambda) \quad \text{for all } \lambda > \overline{\lambda}_\alpha.$$

Thus, the equation $\zeta_\alpha(\lambda_\alpha^*) = \phi(\lambda_\alpha^*)$ becomes $\psi_\alpha(\lambda_\alpha^*) = \phi(\lambda_\alpha^*)$. Using the explicit definitions of these functions given in (13) and (14), we get

$$1/\alpha = \mathbb{E}\frac{zs^2 - z}{\lambda_\alpha^* - z}. \tag{68}$$

We can also explicitly compute

$$\psi'_\alpha(\lambda_\alpha^*) = 1/\alpha - \mathbb{E}\frac{z^2}{(\lambda_\alpha^* - z)^2} \tag{69}$$

$$= \mathbb{E}\frac{zs^2 - z}{\lambda_\alpha^* - z} - \mathbb{E}\frac{z^2}{(\lambda_\alpha^* - z)^2}. \tag{70}$$

Similarly, we can write

$$\phi'(\lambda_\alpha^*) = -\mathbb{E}\frac{z^2 s^2}{(\lambda_\alpha^* - z)^2}. \tag{71}$$

Substituting (70) and (71) into the asymptotic characterization (18) gives us (22), which, together with (68), provides a parametric representation of the function $\rho(\alpha)$.

Finally, we show that $\rho(\alpha) \to 1$ as $\alpha \to \infty$. From (68) and after some simple manipulations, we have

$$\lambda_\alpha^*/\alpha = \mathbb{E}(zs^2 - z) + \mathbb{E}\frac{z^2 s^2 - z^2}{\lambda_\alpha^* - z},$$

Since $\lambda_\alpha^* \to \infty$ as $\alpha \to \infty$, the above formula gives us

$$\lambda_\alpha^* = \mathcal{O}(\alpha),$$

where the leading coefficient $\mathbb{E}(zs^2 - z)$ is positive by assumption (A.6). By the boundedness of z,

$$\frac{\mathbb{E}z^2}{(\lambda_\alpha^*)^2} \leq \mathbb{E}\frac{z^2}{(\lambda_\alpha^* - z)^2} \leq \frac{\mathbb{E}z^2}{(\lambda_\alpha^* - \tau)^2},$$

and thus $\mathbb{E}\frac{z^2}{(\lambda_\alpha^* - z)^2} = \mathcal{O}(1/\alpha^2)$. It follows from (69) that $\psi'_\alpha(\lambda_\alpha^*) = \mathcal{O}(1/\alpha)$. Similarly, we conclude from (71) that $\big|\phi'(\lambda_\alpha^*)\big| = \mathcal{O}(1/\alpha^2)$. Substituting these limiting expressions into (18) then gives us $\lim_{\alpha\to\infty} \rho(\alpha) = 1$, and this completes the proof.

Remark 7: When the set Λ consists of a single element, which is the case for many signal acquisition models we have studied, $\lambda_{c,\min} = \lambda_{c,\max}$ and thus $\alpha_{c,\min} = \alpha_{c,\max}$. There then exists a single critical sampling ratio α_c separating the uncorrelated phase from the correlated one. For $\alpha < \alpha_c$, the estimates from the spectral method is asymptotically orthogonal to $\boldsymbol{\xi}_n$; for $\alpha > \alpha_c$, the estimates will be concentrated on the surface of a right-circular cone whose generating lines make an angle $\theta = \arccos(\sqrt{\rho(\alpha)})$ to the target vector $\boldsymbol{\xi}_n$. The situation is more complicated when Λ contains multiple zero-crossings, in which case a finite number of correlated and uncorrelated phases can alternatively take place between $\alpha_{c,\min}$ and $\alpha_{c,\max}$. A concrete example demonstrating this situation is shown in the next subsection.

C. Multiple Phase Transitions: an Example

Consider the following model:

$$z_i = y_i = \begin{cases} 1, & \text{if } \big|\boldsymbol{a}_i^T \boldsymbol{\xi}_n\big| \in \mathcal{I}_1 \\ \theta, & \text{if } \big|\boldsymbol{a}_i^T \boldsymbol{\xi}_n\big| \in \mathcal{I}_2 \\ 0, & \text{otherwise,} \end{cases}$$

where $0 < \theta < 1$, and $\mathcal{I}_1, \mathcal{I}_2$ are two nonoverlapping intervals on the positive real axis. We also set $\kappa = \|\boldsymbol{\xi}_n\| = 1$, and thus $\boldsymbol{a}_i^T \boldsymbol{\xi}_n$ has the same distribution as a standard normal random variable, denoted by s. Define

$$\beta_\ell = 2\,\mathbb{E}\,\mathbb{1}_{\mathcal{I}_\ell}(s) \quad \text{and} \quad \omega_\ell = 2\,\mathbb{E}\,(s^2 \mathbb{1}_{\mathcal{I}_\ell}(s)),$$

where $\mathbb{1}_{\mathcal{I}_\ell}(\cdot)$ is the indicator function of $\mathcal{I}_\ell$, for $\ell = 1, 2$. As z takes only 3 different values, we can explicitly compute $\Delta(\lambda)$ in (62) as

$$\Delta(\lambda) = \frac{\beta_1 \lambda}{(\lambda - 1)^2} + \frac{\theta\beta_2 \lambda}{(\lambda - \theta)^2} - \frac{\omega_1}{\lambda - 1} - \frac{\theta\omega_2}{\lambda - \theta}.$$

Choose $\theta = 0.48$, $\mathcal{I}_1 = [4.7947, 4.9847]$, and $\mathcal{I}_2 = [0.8995, 0.8998]$. We then have $\beta_1 = 1.0086 \times 10^{-6}$, $\beta_2 = 1.5970 \times 10^{-4}$, $\omega_1 = 2.3976 \times 10^{-5}$ and $\omega_2 = 1.2926 \times 10^{-4}$. In this case, $\Delta(\lambda)$ turns out to have three zero-crossings:

$$\lambda_{c,1} = 1.0765, \quad \lambda_{c,2} = 1.1844, \quad \lambda_{c,3} = 3.3127.$$

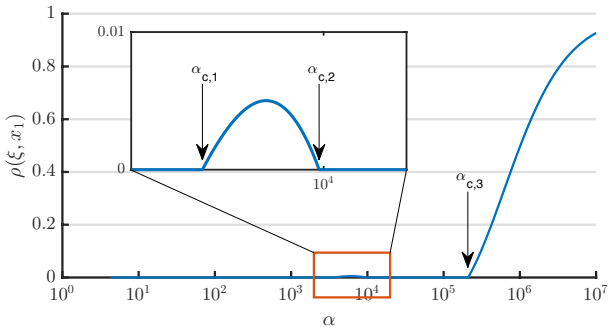

Fig. 4. An example of multiple phase transitions. Shown in the figure is the limiting squared cosine similarity $\rho(\alpha)$ as a function of the sampling ratio α. As α increases, the function $\rho(\alpha)$ alternates between the uncorrelated and correlated phases, with three phase transition points $\{\alpha_{c,i}\}_{1\le i\le 3}$. The insert shows a zoomed-in view of the first correlated phase which takes place within the interval $(\alpha_{c,1}, \alpha_{c,2})$.

By the mapping in (61), they correspond to three critical sampling ratios:

$$\alpha_{c,1} = 3.6279\times 10^3, \alpha_{c,2} = 9.6302\times 10^3, \alpha_{c,3} = 2.0947\times 10^5.$$

Using the characterization given in Theorem 1, we obtain the limiting values of the squared cosine similarity as a function of the sampling ratio α. Figure 4 illustrates this function $\rho(\alpha)$. We can see that, when $\alpha < \alpha_{c,1}$, the estimates given by the spectral method are asymptotically uncorrelated with $\boldsymbol{\xi}_n$. When α is in the interval $(\alpha_{c,1}, \alpha_{c,2})$, however, the function $\rho(\alpha)$ has a small "bump" (see the insert for a zoomed-in view), meaning that the estimates become asymptotically correlated with $\boldsymbol{\xi}_n$. However, the correlation returns to zero as α moves past the second phase transition point $\alpha_{c,2}$. Finally, when $\alpha > \alpha_{c,3}$, the estimates become correlated with $\boldsymbol{\xi}_n$ again, and $\rho(\alpha)$ tends to one as $\alpha \to \infty$.

Remark 8: It would be desirable to obtain a deeper understanding of the above phenomenon involving multiple phase transitions. The example provided here is purely theoretical, as its phase transitions take place at very large values of α. It will be interesting to explore other possible examples of multiple phase transitions with more practical values of α. Moreover, as most signal acquisition models we have studied seem to involve only a single phase transition point, it will be interesting to seek easy-to-verify conditions for the function $\Delta(\lambda)$ defined in (62) to have only one zero-crossing. We leave these as interesting open questions.

V. Discussion

In this paper, we have presented a precise asymptotic characterization of the performance of a spectral method for estimating signals from generalized linear measurements with Gaussian sensing vectors. Our analysis also reveals a phase transition phenomenon that takes place at certain critical sampling ratios. Below a minimum threshold, estimates given by the methods are nearly orthogonal to the true signal $\boldsymbol{\xi}$, thus carrying no information; above a maximum threshold, the estimates become increasingly aligned with $\boldsymbol{\xi}$. The computational complexity of the spectral method is also markedly different in the two phases. Within the uncorrelated phase, the gap between the top two leading eigenvalues diminishes to zero. In contrast, a nonzero spectral gap emerges within the correlated phase. In this section, we close the paper by discussing some possible directions for extending and improving our results as well as their connections to related work in the literature.

The rate of convergence and more refined analysis. The performance of the spectral method was first studied in [11] for the problem of phase retrieval. In that paper, it is shown that, for each $\delta \in (0,1)$, there is a constant $c_1(\delta)$ such that $\rho(\xi_n, x_1^n) > 1-\delta$ with high probability when

$$m > c_1(\delta) n \log^3 n.$$

This estimate of the sample complexity was improved to $m > c_2(\delta) n \log n$ in [7] and to $m > c_3(\delta) n$ in [13]. The key technical tools underlying these previous estimates are matrix concentration inequalities (see, *e.g.*, [37]), which guarantee that the spectral norm of the difference between the data matrix $\boldsymbol{D}_m$ and its expectation $\mathbb{E}\boldsymbol{D}_m$ will be small when the sampling ratio m/n is sufficiently large. The closeness of the corresponding leading eigenvectors of $\boldsymbol{D}_m$ and $\mathbb{E}\boldsymbol{D}_m$ then follow from standard perturbation arguments. (See also our discussions towards the end of Section II-A.) Our work differs from and complements these finite-sample bounds in that we obtain sharp asymptotics to characterize the exact performance of the spectral method in the high-dimensional regime. A (theoretical) limitation of our analysis is that it is asymptotic in nature, requiring both $m, n \to \infty$. Although numerical simulations shown in Section II-D indicate that the asymptotic predictions are accurate even for moderate signal dimensions, it will be useful to quantify the rate of convergence towards the asymptotic limits in future work.

In our analysis, we have assumed that the variable z in (5) is bounded within a finite interval $[0, \tau]$. Whereas z having a finite upper bound is necessary to achieve linear sample complexities (as pointed out in [13]), the assumption that z be nonnegative is largely made to simplify our analysis. We expect that our results can be extended to cases where z is bounded within a finite, but not necessarily nonnegative, interval. Finally, another possible direction to further refine our analysis is to consider second-order asymptotics at the level of central limit theorems (CLTs.) See for instance [38] for a related CLT analysis for the extreme eigenvalues of spiked covariance models.

Alternative initialization schemes. The spectral method considered in this paper is certainly not the only choice for initialization purposes. For example, an interesting alternative is the simple linear estimator studied in [39]:

$$\boldsymbol{x}^n_{\text{linear}} = \frac{1}{m}\sum_{i=1}^m \mathcal{T}(y_i)\boldsymbol{a}_i. \tag{72}$$

By using the moment calculations in [39, Proposition 1.1] and bounding high-order moments, one can easily obtain that

$$\rho(\boldsymbol{\xi}_n, \boldsymbol{x}^n_{\text{linear}}) \xrightarrow{\mathcal{P}} \frac{(\mathbb{E}zs)^2}{(\mathbb{E}zs)^2 + \mathbb{E}z^2/\alpha}, \tag{73}$$

where s and z are the random variables defined in (5).

Recall the function $g(s)$ introduced in (12) and our discussions thereafter, where we point out that the spectral method

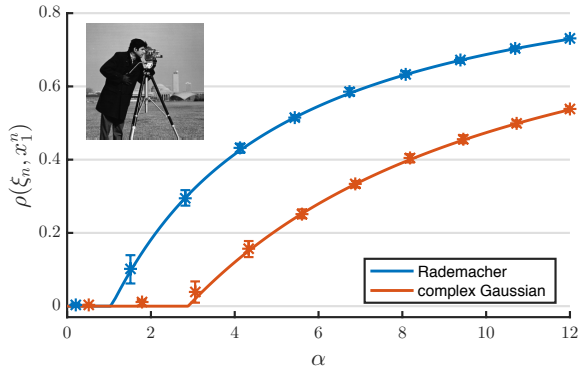

Fig. 5. Demonstrating the potential *universality* of our asymptotic characterizations. The figure shows the results of using the spectral method to estimate a 64×64 *cameraman* image from phaseless measurements under Poisson noise. Theoretical predictions (the blue and red lines) are shown together with simulation results averaged over 16 independent trials. The error bars show one standard deviation. The blue line corresponds to real-valued sensing vectors drawn from the i.i.d. Rademacher ensemble, whereas the red line shows the results of using complex-valued sensing vectors drawn from the complex Gaussian distribution.

is not suitable for acquisition models for which $g(s)$ is an odd function plus a constant. Such cases will pose no problem for the linear estimator in (72). However, it is interesting to note that the linear estimator will be ineffective when $g(s)$ is an even function (as is the case in phase retrieval for which the spectral method was originally developed.) To see this, we note that $\mathbb{E}zs = \mathbb{E}g(s)s = 0$ when $g(s)$ is even. It then follows from (73) that the linear estimator will be asymptotically uncorrelated with the target signal $\boldsymbol{\xi}_n$.

For cases where the function $g(s)$ is neither odd nor even, the choice between the spectral method and the linear estimator is not as clear-cut. The spectral method exhibits phase transition behaviors with its estimates in the uncorrelated phase at small values of α. In contrast, as shown in (73), the performance of the linear estimator increases as a monotonic function of α. As a result, in the regime of very small α, the linear estimator will be preferable. For (moderately) larger values of α, the comparison between the spectral method and the linear estimator cannot be easily made, as their performance also depends on the preprocessing function $\mathcal{T}(\cdot)$ used in (3) and (72). Thus, an interesting avenue for future research is the optimal design of $\mathcal{T}(\cdot)$.

The incorporation of priors. In this work, we assume that the target signal $\boldsymbol{\xi}_n$ is an arbitrary unknown (deterministic) signal. In many applications, the underlying signals satisfy additional constraints (such as sparsity.) In [39], the authors considered a two-step scheme, where the initial linear estimate given in (72) is further projected onto a set which encapsulates one's prior knowledge about $\boldsymbol{\xi}$. It will be interesting to consider and analyze similar projection schemes for the estimates obtained by the spectral method.

Universality and more realistic sensing vectors. Our asymptotic analysis assumes that the sensing vectors are real-valued i.i.d. Gaussian random vectors. Numerical simulations seem to suggest that the theoretical predictions given in Theorem 1 remain valid for more general random measurement ensembles and for complex-valued sensing vectors. To demonstrate this, we show in Figure 5 the results of applying the spectral method to estimate a 64×64 *cameraman* image from phaseless measurements under Poisson noise:

$$y_i \sim \text{Poisson}(|\boldsymbol{a}_i^* \boldsymbol{\xi}|^2) \quad \text{and } z_i = \min\{y_i, \tau\}, \tag{74}$$

where the bound τ is set to 5 and $\|\boldsymbol{\xi}\|$ is normalized to 1 in our simulations. Two measurement ensembles are considered: real-valued sensing vectors whose elements are independent Rademacher (± 1) random variables, and complex-valued sensing vectors with elements drawn from the complex Gaussian distribution $\mathcal{N}(0, \frac{1}{2}) + j\mathcal{N}(0, \frac{1}{2})$. We see from the figure that the theoretical predictions (the solid lines) have excellent agreement with simulation results for this moderately-sized problem, even though the sensing vectors are non-Gaussian. Rigorously establishing the validity of our asymptotic predictions without the Gaussian assumption will be an important future work. A related but more ambitious line of work will be to characterize the performance of the spectral method for structured and more practical sensing ensembles such as the coded diffraction scheme for phase retrieval with random modulation patterns.

APPENDIX

A. Sufficient Conditions for Assumption (A.5) to Hold

In this appendix, we provide two sufficient conditions for Assumption (A.5) to hold.

Case 1: Suppose that the probability law of the random variable z contains a point mass $c\,\delta(z-\tau)$ at its upper boundary τ, where c is some positive constant. This applies to the logistic regression model in Example 1, the subset algorithm (28) in Example 2, the noisy phase retrieval model in (74), and the quantization model described in Section IV-C.

In this case,

$$\mathbb{E}\frac{z}{(\lambda - z)^2} \geq \frac{c\tau}{(\lambda - \tau)^2} \to \infty$$

as $\lambda \to \tau^+$. To verify the second expression in (6), let $h(z) = \mathbb{E}_{s|z}(s^2|z)$. Since $\mathbb{P}(z = \tau) > 0$, we must have $h(\tau) > 0$. Thus,

$$\mathbb{E}\frac{zs^2}{\lambda - z} = \mathbb{E}\frac{zh(z)}{\lambda - z} \geq \frac{c\tau h(\tau)}{\lambda - \tau},$$

which tends to ∞ as λ approaches τ from the right.

Case 2: Suppose that there exist some positive constants c and ε such that the probability density function $p_Z(z)$ of z and the conditional moment $h(z)$ are both bounded below by c for all $z \in [\tau - \varepsilon, \tau]$. The model in (8) represents one such case. Under this setting,

$$\begin{aligned}\mathbb{E}\frac{zs^2}{\lambda - z} &\geq \int_{\tau-\varepsilon}^{\tau} \frac{zh(z)}{\lambda - z} p_Z(z)\,\mathrm{d}z \\ &\geq (\tau - \varepsilon)c^2 \log\left(1 + \frac{\varepsilon}{\lambda - \tau}\right) \xrightarrow{\lambda \to \tau^+} \infty.\end{aligned}$$

Similarly, we can verify that $\mathbb{E}\frac{z}{(\lambda - z)^2} \to \infty$ as $\lambda \to \tau^+$.

B. Norm Estimation

The spectral initialization method estimates the orientation of the vector $\boldsymbol{\xi}_n$ but it provides no information about its norm, as the eigenvector $\boldsymbol{x}_1^n$ is always normalized. In many cases where the sensing vectors come from certain random ensembles, however, the norm $\|\boldsymbol{\xi}_n\|$ can be accurately estimated from the measurements.

As a simple illustrative example, we can consider the (noiseless) phase retrieval problem: $y_i = (\boldsymbol{a}_i^T\boldsymbol{\xi}_n)^2$, where $\boldsymbol{\xi}_n$ is a deterministic unknown vector with $\kappa = \|\boldsymbol{\xi}_n\|$, and the sensing vectors $\{\boldsymbol{a}_i\}$ are i.i.d. standard normal random vectors. Since $\boldsymbol{a}_i^T\boldsymbol{\xi}_n \sim \mathcal{N}(0,\kappa^2)$, the measurement y_i can be represented as

$$y_i \sim \kappa^2 s_i^2,$$

where s_i (for $1 \le i \le m$) are i.i.d. standard normal random variables. A simple estimator of the norm is then

$$\widehat{\kappa} = \sqrt{\frac{\sum_{i=1}^m y_i}{m}}, \tag{75}$$

which is asymptotically consistent as $m \to \infty$.

More generally, consider an observation model $y_i \sim f(y \,|\, \boldsymbol{a}_i^T\boldsymbol{\xi}_n)$, where $f(\cdot\,|\,\cdot)$ is a conditional probability density function and $\boldsymbol{a}_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,\boldsymbol{I}_n)$. Again, writing $\boldsymbol{a}_i^T\boldsymbol{\xi}_n = \kappa s_i$ for i.i.d. normal random variables $\{s_i\}$, we can represent the probability distributions of the measurements $\{y_i\}$ as

$$y_i \overset{\text{i.i.d.}}{\sim} \int_{-\infty}^{\infty} f(y \,|\, \kappa s)\frac{1}{\sqrt{2\pi}}e^{-\frac{s^2}{2}}\,\mathrm{d}s \overset{\text{def}}{=} p_\kappa(y).$$

Let $w(\kappa) \overset{\text{def}}{=} \mathbb{E}(y_i) = \int y\, p_\kappa(y)\,\mathrm{d}y$. If $w(\kappa)$ is monotonic on the positive real line, the *method of moments* gives an estimator

$$\widehat{\kappa}_{\text{MoM}} = w^{-1}\Big(\sum_{i=1}^m y_i/m\Big). \tag{76}$$

We note that the estimator in (75) is a special case of (76). More generally, one could also estimate κ by using maximum likelihood

$$\widehat{\kappa}_{\text{MLE}} = \arg\max_{\kappa>0} \sum_{i\le m} \log\, p_\kappa(y_i),$$

whose asymptotic consistency can be established under standard conditions [40] on the parametric density function $p_\kappa(y)$.

C. Proof of Proposition 2

By a suitable choice of a transformation matrix

$$\widetilde{\boldsymbol{W}} = \left[\begin{array}{c|c} 1 & \boldsymbol{0} \\ \hline \boldsymbol{0} & \boldsymbol{W} \end{array}\right],$$

where $\boldsymbol{W} \in \mathbb{R}^{(n-1)\times(n-1)}$ is an orthogonal matrix involving the last $(n-1)$ rows and columns only, we can get a matrix

$$\widetilde{\boldsymbol{D}} = \widetilde{\boldsymbol{W}}^T\boldsymbol{D}\widetilde{\boldsymbol{W}} = \left[\begin{array}{c|c|c} a & \widetilde{\boldsymbol{q}}^T & \boldsymbol{0} \\ \hline \widetilde{\boldsymbol{q}} & \mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}}\right\}_{i\in\mathcal{I}_1} & \boldsymbol{0} \\ \hline \boldsymbol{0} & \boldsymbol{0} & \mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}}\right\}_{i\in\mathcal{I}_2} \end{array}\right], \tag{77}$$

where $\mathcal{I}_1, \mathcal{I}_2$ are the two sets of indices defined in (42) and $\widetilde{\boldsymbol{q}}$ is a vector consisting of all the nonzero elements of $\boldsymbol{W}^T\boldsymbol{q}$. Let $\lambda_1^{\widetilde{\boldsymbol{D}}}$ and $\widetilde{\boldsymbol{x}}_1$ be the largest eigenvalue of $\widetilde{\boldsymbol{D}}$ and an associated unit-norm eigenvector, respectively. Clearly, $\lambda_1^{\boldsymbol{D}} = \lambda_1^{\widetilde{\boldsymbol{D}}}$ and $(\boldsymbol{e}_1^T\boldsymbol{x}_1)^2 = (\boldsymbol{e}_1^T\widetilde{\boldsymbol{x}}_1)^2$. Thus, we just need to consider $\widetilde{\boldsymbol{D}}$ in our proof.

Due to its block-diagonal form, the eigenvalues of $\widetilde{\boldsymbol{D}}$ is the union of those of its top-left submatrix

$$\boldsymbol{S} = \left[\begin{array}{c|c} a & \widetilde{\boldsymbol{q}}^T \\ \hline \widetilde{\boldsymbol{q}} & \mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}}\right\}_{i\in\mathcal{I}_1} \end{array}\right].$$

and those of its bottom-right submatrix $\mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$. In particular,

$$\lambda_1^{\widetilde{D}} = \lambda_1^{\boldsymbol{S}} \vee \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}. \tag{78}$$

The eigenvectors associated with $\mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$ are easy to characterize. Clearly, each $\lambda_i^{\boldsymbol{P}}, i \in \mathcal{I}_2$ is an eigenvalue of $\widetilde{\boldsymbol{D}}$, and it corresponds to an eigenvector $\boldsymbol{e}_{j(i)}$, where $j(i) \ge 3$ is the row index of $\lambda_i^{\boldsymbol{P}}$ in $\widetilde{\boldsymbol{D}}$.

The eigenvalues and eigenvectors of $\boldsymbol{S}$ can also be precisely characterized. Due to its shape, $\boldsymbol{S}$ is sometimes referred to in the literature as an arrowhead matrix [41], [42]. It can be shown (see for instance [43][pp. 94 – 97]) that $\lambda_1^{\boldsymbol{S}}$ is the unique point within the interval $\lambda > \max\{\lambda_i : i \in \mathcal{I}_1\}$ to satisfy the equation

$$a = \lambda_1^{\boldsymbol{S}} + R(\lambda_1^{\boldsymbol{S}}), \tag{79}$$

where $R(\lambda)$ is the function defined in (38). (Alternatively, we can use the Laplace expansion to explicitly derive the characteristic polynomial of $\boldsymbol{S}$ as

$$(\lambda - a)\prod_{i\in\mathcal{I}_1}(\lambda - \lambda_i^{\boldsymbol{P}}) - \sum_{i\in\mathcal{I}_1}(\boldsymbol{w}_i^T\boldsymbol{q})^2 \prod_{j\in\mathcal{I}_1\setminus i}(\lambda - \lambda_j^{\boldsymbol{P}}).$$

Then, by following similar arguments as those used in the proof of Lemma 1, we can reach the characterization (79) about $\lambda_1^{\boldsymbol{S}}$.) Furthermore, let $\boldsymbol{x}_1^{\boldsymbol{S}}$ be a unit-norm eigenvector of $\widetilde{\boldsymbol{D}}$ associated with $\lambda_1^{\boldsymbol{S}}$. It is easily checked that

$$\boldsymbol{x}_1^{\boldsymbol{S}} = \begin{bmatrix} 1 & \boldsymbol{y} & \boldsymbol{0}_r \end{bmatrix}/(1 + \|\boldsymbol{y}\|^2)^{1/2}, \tag{80}$$

where $\boldsymbol{y} = (\lambda_1^{\boldsymbol{S}}\boldsymbol{I} - \mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}}\right\}_{i\in\mathcal{I}_1})^{-1}\widehat{\boldsymbol{q}}$ and $\boldsymbol{0}_r$ is a row vector of r zeroes with r being the cardinality of $\mathcal{I}_2$. It follows that

$$\begin{aligned}(\boldsymbol{e}_1^T\boldsymbol{x}_1^{\boldsymbol{S}})^2 &= \Big(1 + \widetilde{\boldsymbol{q}}^T\big(\lambda_1^{\boldsymbol{S}}\boldsymbol{I} - \mathrm{diag}\left\{\lambda_i^{\boldsymbol{P}}\right\}_{i\in\mathcal{I}_1}\big)^{-2}\widetilde{\boldsymbol{q}}\Big)^{-1} \\ &= \big(1 + \boldsymbol{q}^T(\lambda_1^{\boldsymbol{S}}\boldsymbol{I} - \boldsymbol{P})^{-2}\boldsymbol{q}\big)^{-1} \\ &= \big(1 + R'(\lambda_1^{\boldsymbol{S}})\big)^{-1}, \end{aligned} \tag{81}$$

where $R'(\lambda)$ denotes the derivative of the function $R(\lambda)$.

To show the claim of the proposition, we consider the following three cases.

Case 1: $\lambda_1^{\boldsymbol{S}} > \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$. We choose $\mu^* = -1/R(\lambda_1^{\boldsymbol{S}})$, and thus $\lambda_1^{\boldsymbol{S}} = R^{-1}(-1/\mu^*)$. It follows from Lemma 1 that

$$L(\mu^*) = \lambda_1^{\boldsymbol{S}} \vee \lambda_1^{\boldsymbol{P}} = \lambda_1^{\boldsymbol{S}} = \lambda_1^{\widetilde{\boldsymbol{D}}},$$

where the second equality is due to the fact that

$$\lambda_1^{\boldsymbol{S}} > \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_1\right\}, \tag{82}$$

and the last equality comes from (78).

Using the identity (79) for $\lambda_1^{\boldsymbol{S}}$, we can also verify that μ^* indeed satisfies the equation (44). (Its uniqueness is always guaranteed; see Remark 5 at the end of Section III-B.) The unit-norm leading eigenvector of $\widetilde{\boldsymbol{D}}$ in this case is the vector $\boldsymbol{x}_1^{\boldsymbol{S}}$ defined in (80). Since $L(\mu) = R^{-1}(-1/\mu)$ in a neighborhood of μ^*, the function $L(\mu)$ is differentiable at μ^* and

$$L'(\mu^*) = \left(1/R'(\lambda_1^{\boldsymbol{S}})\right)(\mu^*)^{-2}. \tag{83}$$

Substituting (83) into (81) leads to (46).

Case 2: $\lambda_1^{\boldsymbol{S}} < \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$, in which case $\lambda_1^{\widetilde{D}} = \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\} = \lambda_1^{\boldsymbol{P}}$, where the last equality is due to (82). The corresponding leading eigenvector has nonzero elements only in its last r entries, where r is the cardinality of $\mathcal{I}_2$. Thus,

$$(\boldsymbol{e}_1^T \widetilde{\boldsymbol{x}}_1)^2 = 0. \tag{84}$$

We set $\mu^* = (\lambda_1^{\boldsymbol{P}} - a)^{-1}$. (Note that we are guaranteed to have $\mu^* > 0$. This can be verified by observing that $\lambda_1^{\boldsymbol{P}} > \lambda_1^{\boldsymbol{S}} = a - R(\lambda_1^{\boldsymbol{S}}) > a$, where the equality is due to (79) and the last inequality follows from the fact that $R(\lambda) < 0$.) Since $R(\lambda)$ is a strictly increasing function, we have

$$R^{-1}(-1/\mu^*) < R^{-1}(a - \lambda_1^{\boldsymbol{S}}) = \lambda_1^{\boldsymbol{S}} < \lambda_1^{\boldsymbol{P}},$$

where the equality comes from (79). It then follows from Lemma 1 that $L(\mu^*) = \lambda_1^{\boldsymbol{P}} = \lambda_1^{\widetilde{\boldsymbol{D}}}$ and, moreover, μ^* satisfies the equation (44).

To characterize the eigenvector, we note that $L(\mu) \equiv \lambda_1^{\boldsymbol{P}}$ in a neighborhood of μ^*. We then have $L'(\mu^*) = 0$, which, together with (84), leads to (46).

Case 3: $\lambda_1^{\boldsymbol{S}} = \max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$. This is a special case, where the algebraic multiplicity of the leading eigenvalue $\lambda_1^{\widetilde{\boldsymbol{D}}} = \lambda_1^{\boldsymbol{S}} = \lambda_1^{\boldsymbol{P}}$ is greater than one. The leading eigenvectors are not unique, and they can be any vector in the form of

$$c_1 \boldsymbol{x}_1^{\boldsymbol{S}} + c_2 \boldsymbol{v},$$

where $\boldsymbol{x}_1^{\boldsymbol{S}}$ is the eigenvector defined in (80) and $\boldsymbol{v}$ is an eigenvector associated with $\max\left\{\lambda_i^{\boldsymbol{P}} : i \in \mathcal{I}_2\right\}$, and c_1, c_2 are two constants satisfying $c_1^2 + c_2^2 = 1$. Since $\boldsymbol{e}_1^T \boldsymbol{v} = 0$, we have from (81) that

$$(\boldsymbol{e}_1^T \widetilde{\boldsymbol{x}}_1)^2 \in [0, \left(1 + R'(\lambda_1^{\boldsymbol{S}})\right)^{-1}]. \tag{85}$$

Same as what we did in Case 2, we set $\mu^* = (\lambda_1^{\boldsymbol{P}} - a)^{-1}$. Following the same arguments there, we can show that $L(\mu^*) = \lambda_1^{\boldsymbol{P}} = \lambda_1^{\widetilde{\boldsymbol{D}}}$ and μ^* satisfies the equation (44). Moreover, we can see that $L(\mu) = R^{-1}(-1/\mu)$ for $\mu > \mu^*$ and $L(\mu) \equiv \lambda_1^{\boldsymbol{P}}$ for $\mu < \mu^*$. The function $L(\mu)$ is not differentiable at μ^*, but its right and left derivatives do exist. It is easy to get $\partial_+ L(\mu^*) = \left(1/R'(\lambda_1^{\boldsymbol{S}})\right)(\mu^*)^{-2}$ [see (83)] and $\partial_- L(\mu^*) = 0$. Substituting these quantities into (85), we reach the characterization given in (45).

D. Proof of Proposition 3

To establish the almost-sure convergence of the random measure $f^{\boldsymbol{M}_m}(\lambda)$ to the probability law of z, we just need to show that, almost surely, the empirical distribution function

$$F^{\boldsymbol{M}_m}(\lambda) = \frac{1}{m-1} \# \left\{2 \le j \le m : \lambda_j^{\boldsymbol{M}_m} \le \lambda\right\}$$

converges to $F_z(\lambda)$, the cumulative distribution function of z, at all points λ where $F_z(\lambda)$ is continuous. Since $\boldsymbol{M}_m$ is a rank-one perturbation of the diagonal matrix $\boldsymbol{Z}$, standard interlacing theorems (see [35, Theorem 4.3.4]) give us

$$\lambda_k^{\boldsymbol{Z}} \ge \lambda_{k+1}^{\boldsymbol{M}_m} \ge \lambda_{k+2}^{\boldsymbol{Z}}, \tag{86}$$

for $1 \le k \le m-2$. Let $F^{\boldsymbol{Z}}(\lambda) = \frac{1}{m} \# \left\{1 \le j \le m : z_j \le \lambda\right\}$ be the empirical distribution function of the eigenvalues of $\boldsymbol{Z}$. We can then easily verify from (86) that

$$mF^{\boldsymbol{Z}}(\lambda) - 2 \le (m-1)F^{\boldsymbol{M}_m}(\lambda) \le mF^{\boldsymbol{Z}}(\lambda) + 1. \tag{87}$$

Since $\{z_i\}_{1 \le i \le m}$ is an i.i.d. sample of the random variable z, with probability one $F^{\boldsymbol{Z}}(\lambda)$ converges to $F_z(\lambda)$ all all points λ where $F_z(\lambda)$ is continuous. It then follows from (87) that $F^{\boldsymbol{M}_m}(\lambda)$ converges almost surely to the same limit $F_z(\lambda)$.

To study the leading eigenvalue $\lambda_1^{\boldsymbol{M}_m}$, we use Lemma 1. To apply that result, we require $\boldsymbol{v} = [z_1 s_1, z_2 s_2, \ldots, z_m s_m]^T$ as defined in (35) to be not equal to the all-zero vector. This condition holds almost surely for all sufficiently large m. To see this, we note that $\mathbb{P}(z_i = 0) < 1$, as otherwise assumption (A.6) will not hold. Moreover, $s_i \neq 0$ with probability one. It follows that the i.i.d. sequence $z_1 s_1, z_2 s_2, z_3 s_3, \ldots$ has an infinite number of nonzero elements. Thus, almost surely, the m-dimensional vector $\boldsymbol{v} \neq \boldsymbol{0}$ for sufficiently large m.

Applying (39) to our case, we have $\lambda_1^{\boldsymbol{M}_m} = R_m^{-1}(-1/\mu) \vee \max\{z_i\}_{1 \le i \le m}$, where

$$R_m(\lambda) = \sum_{i=1}^{m} \frac{z_i^2 s_i^2}{z_i - \lambda}$$

with this function defined on $\lambda > \max\{z_i\}_{1 \le i \le m}$. Since $R_m^{-1}(-1/\mu) > \max\{z_i\}_{1 \le i \le m}$, we can further simplify the characterization to

$$\lambda_1^{\boldsymbol{M}_m} = R_m^{-1}(-1/\mu).$$

For every $\lambda > \tau$, with τ being the upper bound of the support of the probability distribution of z, it follows from the strong law of large numbers that $R_m(\lambda)$ converges almost surely to

$$\mathbb{E}\frac{z^2 s^2}{z - \lambda} = -Q(\lambda),$$

where $Q(\lambda)$ is defined in (51). On its domain $\lambda > \tau$, the function $-Q(\lambda)$ is strictly increasing and thus it admits a functional inverse $(-Q)^{-1}(x) = Q^{-1}(-x)$. Applying Lemma 3 in Appendix E, we have

$$\lambda_1^{\boldsymbol{M}_m} = R_m^{-1}(-1/\mu) \xrightarrow{\text{a.s.}} Q^{-1}(1/\mu)$$

as $m \to \infty$.

E. Auxiliary Lemmas

We prove here two auxiliary lemmas that are used in our proofs of Proposition 3 and Theorem 1.

Lemma 3: Let $\{f_n(x)\}_{n\geq 1}$ be a family of (random) functions defined on an open interval (a,b). Each $f_n(x)$ is continuous and nondecreasing. For each $x\in(a,b)$, $f_n(x)\xrightarrow{\text{a.s.}} f(x)$ as $n\to\infty$, where $f(x)$ is a continuous and nondecreasing function. Then, for any sequence $\{x_n\}\subset(a,b)$ with $x_n\xrightarrow{\text{a.s.}} x^*\in(a,b)$, we have

$$f_n(x_n)\xrightarrow{\text{a.s.}} f(x^*). \tag{88}$$

If, in addition, the functions $\{f_n(x)\}$ and $f(x)$ are strictly increasing, we denote by $\{f_n^{-1}(x)\}_{n\geq 1}$ and $f^{-1}(x)$ the corresponding functional inverses. Assume that the domains of $\{f_n^{-1}(x)\}_{n\geq 1}$ and $f^{-1}(x)$ contain a common open interval $\mathcal{I}$. Then for any sequence $\{y_n\}_{n\geq 1}\subset\mathcal{I}$ such that $y_n\xrightarrow{\text{a.s.}} y\in\mathcal{I}$, we have

$$f_n^{-1}(y_n)\xrightarrow{\text{a.s.}} f^{-1}(y). \tag{89}$$

Proof: We first show (88). Let $\beta_k=x^*-h/k$ for $k=1,2,\ldots$ be a sequence that converges to x^* from the left. We choose $h<\min\{x^*-a,b-x^*\}$ so that the entire sequence stays within the interval (a,b). Similarly, define a sequence $\gamma_k=x^*+h/k$, for $k=1,2,\ldots$, that converges to x^* from the right. Denote by $\mathcal{A}$ the intersection of the event that $f_n(x)\to f(x)$ for all $x\in\{\beta_k\}_k\cup\{\gamma_k\}_k$ and the event that $x_n\to x$. Clearly, $\mathbb{P}(\mathcal{A})=1$. Next, we show that (88) holds within this almost sure event.

Fix $k\geq 1$. As $x_n\to x^*$, we have $\beta_k\leq x_n\leq\gamma_k$ for all sufficiently large n. By the monotonicity of $f_n(x)$,

$$f_n(\beta_k)\leq f_n(x_n)\leq f_n(\gamma_k).$$

It follows that

$$f(\beta_k)\leq\liminf_n f_n(x_n)\leq\limsup_n f_n(x_n)\leq f(\gamma_k).$$

As k is arbitrary, we take the $k\to\infty$ limit, which leads to $\lim_n f_n(x_n)=f(x)$ by the continuity of $f(x)$.

The proof of (89) is similar. We establish it under the additional assumption that $\{f_n(x)\}$ and $f(x)$ are strictly increasing. Construct two sequences $\{\beta_k\}$ and $\{\gamma_k\}$ as above, with x^* replaced by $f^{-1}(y)$. Also define the event $\mathcal{A}$ similarly. We show that, within the almost sure event $\mathcal{A}$, we have $f_n^{-1}(y_n)\to f^{-1}(y)$.

Fix $k\geq 1$. Since $f(x)$ is strictly increasing, $\beta_k<f^{-1}(y)<\gamma_k$ implies that

$$f(\beta_k)<y<f(\gamma_k).$$

As $f_n(\beta_k)\to f(\beta_k)$, $f_n(\gamma_k)\to f(\gamma_k)$ and $y_n\to y$, the inequalities

$$f_n(\beta_k)<y_n<f_n(\gamma_k),$$

hold for all sufficiently large n. By the strict monotonicity of $f_n(x)$,

$$\beta_k<f_n^{-1}(y_n)<\gamma_k,$$

for all sufficiently large n. It then follows that $\beta_k\leq\liminf_n f_n^{-1}(y_n)\leq\limsup_n f_n^{-1}(y_n)\leq\gamma_k$, for each k. As $\beta_k\to f^{-1}(y)$ and $\gamma_k\to f^{-1}(y)$, we are done. ∎

Lemma 4: Let $\{f_n(x)\}_{n\geq 1}$ be a sequence of (random) convex functions defined on an open interval (a,b). For each $x\in(a,b)$, $f_n(x)\xrightarrow{\text{a.s.}} f(x)$. Let $\{x_n\}_{n\geq 1}\subset(a,b)$ be a sequence such that $x_n\xrightarrow{\text{a.s.}} x^*$ for some $x^*\in(a,b)$. If $f(x)$ is differentiable at x^*, then

$$\partial_- f_n(x_n)\xrightarrow{\text{a.s.}} f'(x^*)\quad\text{and}\quad\partial_+ f_n(x_n)\xrightarrow{\text{a.s.}} f'(x^*), \tag{90}$$

where $\partial_- f_n(x)$ and $\partial_+ f_n(x)$ denote the left and right derivatives of $f_n(x)$, respectively.

Proof: Similar to the proof of Lemma 3, we construct two sequences: $\{\beta_k\}_{k\geq 1}$ is strictly increasing and converges to x^* from the left, whereas $\{\gamma_k\}_{k\geq 1}$ is strictly decreasing and converges to x^* from the right. Denote by $\mathcal{A}$ the intersection of the event that $f_n(x)\to f(x)$ for all $x\in\{\beta_k\}_k\cup\{\gamma_k\}_k$ and the event that $x_n\to x^*$. It is easily checked that $\mathbb{P}(\mathcal{A})=1$. Next, we establish (90) within this almost sure event.

For any $i<j$, since $\beta_i<\beta_j<x^*$ and $x_n\to x^*$, we must have $\beta_i<\beta_j<x_n$ for all sufficiently large n. By the convexity of $f_n(x)$, its left derivatives always exist and we have

$$\partial_- f_n(x_n)\geq\frac{f_n(\beta_i)-f_n(\beta_j)}{\beta_i-\beta_j},$$

for all sufficiently large n. It follows that

$$\liminf_n\partial_- f_n(x_n)\geq\frac{f(\beta_i)-f(\beta_j)}{\beta_i-\beta_j}$$

for all $i<j$. Since

$$\lim_{i\to\infty}\left(\lim_{j\to\infty}\frac{f(a_i)-f(a_j)}{a_i-a_j}\right)=f'(x^*),$$

we must have

$$\liminf_n\partial_- f_n(x_n)\geq f'(x^*). \tag{91}$$

Working with the sequence $\{\gamma_k\}_{k\geq 1}$ and using similar arguments as above, we can show that

$$\limsup_n\partial_+ f_n(x_n)\leq f'(x^*). \tag{92}$$

Since $\limsup_n\partial_- f_n(x_n)\leq\limsup_n\partial_+ f_n(x_n)$, we use (91) and (92) to conclude that $\lim_n\partial_- f_n(x_n)$ exists and that it is equal to $f'(x^*)$. By similar arguments, the same claim also holds for the sequence $\{\partial_+ f_n(x_n)\}$, and thus the proof is complete. ∎

REFERENCES

[1] M. Unser and M. Eden, "Maximum likelihood estimation of linear signal parameters for Poisson processes," *IEEE Trans. Acoust., Speech, and Signal Process.*, vol. 36, no. 6, pp. 942–945, Jun. 1988.

[2] F. Yang, Y. M. Lu, L. Sbaiz, and M. Vetterli, "Bits from photons: Oversampled image acquisition using binary poisson statistics," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1421–1436, 2012.

[3] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Applied Optics*, vol. 21, no. 15, pp. 2758–2769, 1982.

[4] S. Rangan and V. K. Goyal, "Recursive consistent estimation with bounded noise," *Information Theory, IEEE Transactions on*, vol. 47, no. 1, pp. 457–464, 2001.

[5] A. J. Dobson and A. Barnett, *An Introduction to Generalized Linear Models*, 3rd ed. Boca Raton: Chapman and Hall/CRC, May 2008.

[6] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*, 2nd ed. New York, NY: Springer, Apr. 2011.

[7] E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.

[8] E. J. Candes and X. Li, "Solving quadratic equations via PhaseLift when there are about as many equations as unknowns," *Foundations of Computational Mathematics*, vol. 14, no. 5, pp. 1017–1026, 2014.

[9] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Convex algorithms and limitations," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 1022–1026.

[10] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Mathematical Programming*, vol. 149, no. 1-2, pp. 47–81, 2015.

[11] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," in *Advances in Neural Information Processing Systems*, 2013, pp. 2796–2804.

[12] E. J. Candes, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *Information Theory, IEEE Transactions on*, vol. 61, no. 4, pp. 1985–2007, 2015.

[13] Y. Chen and E. J. Candes, "Solving Random Quadratic Systems of Equations Is Nearly as Easy as Solving Linear Systems," *arXiv preprint arXiv:1505.05114*, 2015.

[14] H. Zhang and Y. Liang, "Reshaped Wirtinger Flow for Solving Quadratic System of Equations," in *Advances in Neural Information Processing Systems*, 2016, pp. 2622–2630.

[15] G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving Systems of Random Quadratic Equations via Truncated Amplitude Flow," *arXiv:1605.08285*, May 2016.

[16] M. Soltanolkotabi, "Structured signal recovery from quadratic measurements: Breaking sample complexity barriers via nonconvex optimization," *arXiv:1702.06175 [cs, math, stat]*, Feb. 2017.

[17] G. Li, Y. Gu, and Y. M. Lu, "Phase retrieval using iterative projections: Dynamics in the large systems limit," in *Proc. Allerton Conference on Communication, Control and Computing.*, Monticello, IL, Oct 2015.

[18] Y. Chi and Y. M. Lu, "Kaczmarz method for solving quadratic equations," *IEEE Signal Process. Lett.*, vol. 23, no. 9, 2016.

[19] T. Goldstein and C. Studer, "PhaseMax: Convex Phase Retrieval via Basis Pursuit," *arXiv:1610.07531 [cs, math]*, Oct. 2016.

[20] S. Bahmani and J. Romberg, "Phase Retrieval Meets Statistical Learning Theory: A Flexible Convex Relaxation," *arXiv:1610.04210 [cs, math, stat]*, Oct. 2016.

[21] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2980–2998, Jun. 2010.

[22] P. Jain, P. Netrapalli, and S. Sanghavi, "Low-rank matrix completion using alternating minimization," in *Proc. Symposium on Theory of Computing*, 2013.

[23] K. Lee, Y. Li, M. Junge, and Y. Bresler, "Blind recovery of sparse signals from subsampled convolution," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 802–821, Feb. 2017.

[24] X. Li, S. Ling, T. Strohmer, and K. Wei, "Rapid, robust, and reliable blind rapid, robust, and reliable blind deconvolution via nonconvex optimization," *arXiv:1606.04933 [cs.IT]*, 2016.

[25] S. Arora, R. Ge, T. Ma, and A. Moitra, "Simple, efficient, and neural algorithms for sparse coding." in *COLT*, 2015, pp. 113–149.

[26] Y. Chen and E. Candès, "The projected power method: An efficient algorithm for joint alignment from pairwise differences," *arXiv:1609.05820 [cs.IT]*, 2016.

[27] I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *The Annals of Statistics*, vol. 29, no. 2, pp. 295–327, Apr. 2001.

[28] J. Baik, G. B. Arous, and S. Péché, "Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices," *The Annals of Probability*, vol. 33, no. 5, pp. 1643–1697, Sep. 2005.

[29] Z. D. Bai and J. Yao, "On sample eigenvalues in a generalized spiked population model," *Journal of Multivariate Analysis*, vol. 106, no. 167–177, 2012.

[30] Z. Bai and J. W. Silverstein, *Spectral Analysis of Large Dimensional Random Matrices*, ser. Springer Series in Statistics. New York, NY: Springer New York, 2010.

[31] V. A. Marčenko and L. A. Pastur, "Distribution of eigenvalues for some sets of random matrices," *Mathematics of the USSR-Sbornik*, vol. 1, no. 4, p. 457, 1967.

[32] J. W. Silverstein and Z. D. Bai, "On the empirical distribution of eigenvalues of a class of large dimensional random matrices," *Journal of Multivariate analysis*, vol. 54, no. 2, pp. 175–192, 1995.

[33] F. Benaych-Georges and R. R. Nadakuditi, "The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices," *Advances in Mathematics*, vol. 227, no. 1, pp. 494–521, May 2011.

[34] D. A. Harville, *Matrix Algebra from a Statistician's Perspective*. Springer, 2000.

[35] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1985.

[36] J. W. Silverstein and S. I. Choi, "Analysis of the Limiting Spectral Distribution of Large Dimensional Random Matrices," *Journal of Multivariate Analysis*, vol. 54, no. 2, pp. 295–309, Aug. 1995.

[37] R. Vershynin, "Introduction to the non-asymptotic analysis of random matrices," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge, UK: Cambridge University Press, 2012, pp. 210 – 268.

[38] Z. Bai and J.-F. Yao, "Central limit theorems for eigenvalues in a spiked population model," *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, vol. 44, no. 3, pp. 447–474, Jun. 2008.

[39] Y. Plan, R. Vershynin, and E. Yudovina, "High-dimensional estimation with geometric constraints," *Information and Inference: A Journal of the IMA*, vol. 6, no. 1, pp. 1–40, Mar. 2017.

[40] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. New York: Springer, Sep. 2003.

[41] D. P. O'Leary and G. W. Stewart, "Computing the eigenvalues and eigenvectors of symmetric arrowhead matrices," *Journal of Computational Physics*, vol. 90, no. 2, pp. 497–505, Oct. 1990.

[42] N. J. Stor, I. Slapnicar, and J. L. Barlow, "Accurate eigenvalue decomposition of arrowhead matrices and applications," *Linear Algebra and its Applications*, vol. 464, pp. 62–89, Jan. 2015.

[43] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*. Oxford, UK: Clarendon Press, Apr. 1988.